



1

**Interim Report:
Distributed Problem Solving: Adaptive
Networks With a Computer Intermediary
Resource: Intelligent Executive
Computer Communication**

John Lyman and Carla J. Conaway

University of California at Los Angeles

for

**Contracting Officer's Representative
Michael Drillings**

**Office of Basic Research
Michael Kaplan, Director**

June 1991



91-15016



**United States Army Research Institute
for the Behavioral and Social Sciences**

Approved for public release, distribution is unlimited.

91 1104 095

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS --		
2a. SECURITY CLASSIFICATION AUTHORITY --			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE --			5. MONITORING ORGANIZATION REPORT NUMBER(S) ARI Research Note 91-80		
4. PERFORMING ORGANIZATION REPORT NUMBER(S) --			7a. NAME OF MONITORING ORGANIZATION U.S. Army Research Institute		
6a. NAME OF PERFORMING ORGANIZATION University of California at Los Angeles		6b. OFFICE SYMBOL (If applicable) --		7b. ADDRESS (City, State, and ZIP Code) 5001 Eisenhower Avenue Alexandria, VA 22333-5600	
6c. ADDRESS (City, State, and ZIP Code) Dept. of Materials Science and Engineering School of Engineering and Applied Sciences UCLA, Los Angeles, CA 90024		8a. NAME OF FUNDING/SPONSORING ORGANIZATION U.S. Army Research Institute for the Behavioral and Social Sciences		8b. OFFICE SYMBOL (If applicable) PERI-BR	
8c. ADDRESS (City, State, and ZIP Code) 5001 Eisenhower Avenue Alexandria, VA 22333-5600		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER MDA903-84-C-0355			
		10. SOURCE OF FUNDING NUMBERS			
		PROGRAM ELEMENT NO. 61102B	PROJECT NO. 74F	TASK NO. N/A	WORK UNIT ACCESSION NO. N/A
11. TITLE (Include Security Classification) Interim Report: Distributed Problem Solving: Adaptive Networks With a Computer Intermediary Resource: Intelligent Executive Computer Communication					
12. PERSONAL AUTHOR(S) Lyman, John; and Conaway, Carla J.					
13a. TYPE OF REPORT Interim		13b. TIME COVERED FROM 84/10 TO 88/06		14. DATE OF REPORT (Year, Month, Day) 1991, June	
15. PAGE COUNT 73					
16. SUPPLEMENTARY NOTATION Contracting Officer's Representative, Michael Drillings.					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP	Intelligent systems		
			Problem solving		
19. ABSTRACT (Continue on reverse if necessary and identify by block number) This report presents a system concept for an "intelligent executive." Specifically, the functional requirements of a distributed group of agents are analyzed, along with the characteristics of different forms of communication. The intelligent executive embodies guidelines for specifying an appropriate communication strategy given the problem environment and the informational requirements of the group. When employed within a distributed problem solving environment, the intelligent executive is expected to facilitate the solution to a problem by increasing communication efficiency and problem-solving capabilities of distributed agents.					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a. NAME OF RESPONSIBLE INDIVIDUAL Michael Drillings			22b. TELEPHONE (Include Area Code) (202) 274-8722		22c. OFFICE SYMBOL PERI-BR

INTERIM REPORT: DISTRIBUTED PROBLEM SOLVING: ADAPTIVE NETWORKS WITH A
COMPUTER INTERMEDIARY RESOURCE: INTELLIGENT EXECUTIVE COMPUTER COMMUNICATION

CONTENTS

	Page
INTRODUCTION	1
Overview	1
Characteristics of Distributed Problem Solving	2
Challenges in Distributed Problem Solving	6
LITERATURE REVIEW	9
Major Research Issues	9
Current Research Emphasis	17
INTELLIGENT EXECUTIVE FUNCTIONAL REQUIREMENTS	20
Overview	20
Communication Protocols	21
Command Situations	25
Communication Factors	27
Knowledge Requirements	32
INTELLIGENT EXECUTIVE CONCEPT OF OPERATION	39
Overview	39
Conceptual Framework	39
Concept of Operation	45
Selection Environment	48
CONCLUSIONS	54
Overview	54
Experimental Objective	55
Experimental Set-Up	55
Experimental Variables	56
Summary	58
BIBLIOGRAPHY	59

CONTENTS (Continued)

Page

LIST OF TABLES

Table 1.	Communication protocols	22
2.	Command situations	26
3.	Intelligent executive domain and control knowledge	38
4.	Costs per command situation	42

LIST OF FIGURES

Figure 1.	Distributed problem-solving situation	3
2.	Cost factors	31
3.	Risk factors	33
4.	Knowledge base elements to be shared among agents	34
5.	Intelligent executive interpretation of domain knowledge	36
6.	Ratio of shared information	44
7.	Intelligent executive concept of operation	46
8.	Intelligent executive computational results	49
9.	External constraints on expected cost of detection	51
10.	External constraint on expected cost of error	52

**INTERIM REPORT: DISTRIBUTED PROBLEM SOLVING:
ADAPTIVE NETWORKS WITH A COMPUTER INTERMEDIARY RESOURCE:
INTELLIGENT EXECUTIVE COMPUTER COMMUNICATION**

1.0 INTRODUCTION

1.1 OVERVIEW

Problem solving situations are integral to our daily existence. Individuals must recognize, analyze and solve problems in order to achieve goals that they set for themselves or that others set for them. In each case, individuals may have access to many strategies which they can employ to reach a solution. When groups of individuals engage in solving a common problem, the problem solving situation becomes significantly more complex. Strategies that groups can employ to solve problems may not be clear to the members of the group. If the group is not physically colocated (i.e., members are distributed), problem solving becomes even more complex because, not only is the use of available strategies unclear, but also their utility may be doubtful. The purpose of this report is to analyze the requirements for problem solving strategies for the distributed problem solving situation.

In recent years there has been a growing interest in distributed problem solving, i.e., collections of intelligent agents (humans or intelligent machines) which cooperate to solve problems. The motivation for research in this area is to increase the efficiency and capabilities of groups which must solve problems in distributed information situations. To achieve these goals, strategies for distributed problem solving must be developed and empirically tested. Some of these methods may come in the form of "meta-level" functions that are embedded within the group structure. Meta-level functions could typically include providing the agents with high-level knowledge about the capabilities of other agents or providing guidelines on how and when to use the available resources (communication, sensing, computation, etc.).

This report presents a system concept for an "Intelligent Executive." Specifically, the functional requirements of a distributed group of agents are analyzed along with the

characteristics of different forms of communication. The Intelligent Executive embodies guidelines for specifying an appropriate communication strategy given the problem environment and the informational requirements of the group. When employed within a distributed problem solving environment, the Intelligent Executive can be expected to facilitate the solution to the problem by increasing communication efficiency and problem solving capabilities of the distributed agents.

1.2 CHARACTERISTICS OF DISTRIBUTED PROBLEM SOLVING

Distributed problem solving refers to the process by which several agents, spatially separated, interact to solve a problem and achieve a common goal. A familiar metaphor for the distributed problem solving process is a group of human experts (each with individual skills, specialties, and bodies of knowledge) who are capable of sophisticated problem solving and who cooperate with each other to solve a problem. A representation for a distributed problem solving situation is shown in Figure 1. The ellipses depicting **bodies of knowledge** (A,B, and C) each represent the combined skills, knowledge, and capabilities of a problem solving agent, either human or intelligent machine. The **shared knowledge**, R, represents the elements of skills, knowledge, and/or capabilities common to all agents within the problem solving group. It is this commonality that allows the agents to functionally communicate and cooperate with each other effectively. Consequently, one overall goal of the problem solvers is to increase the amount of shared and common knowledge through communication until it encompasses the **solution requirements** of the problem. The solution requirements, R', represent the total resultant shared knowledge that must be possessed by the agents in order to achieve a solution to the problem. It is the minimum amount of commonality required of their bodies of knowledge to be able to reach a solution to the problem at hand.




Figure 1.
Distributed Problem-Solving Situation

Several general characteristics of the distributed problem solving situation are particularly important to understanding the difficulties associated with this type of process. They may also affect the efficiency and productivity of the problem solving process. These characteristics are discussed below.

1.2.1 Skills: Various skills such as sensing, communicating, planning, and acting are brought to bear on the problem by a collection of agents.

The agents may possess differing levels of proficiency in all of these areas.

1.2.2 Tasks: Unlike the situation with single-agent problem solving in which an agent has a global view of the problem, tasks must somehow be assigned to appropriate agents by the group even when no agent has a global view of the situation. Tasks must also be decomposed into, sometimes logically dependent, subtasks.

1.2.3 Knowledge: An agent may be subject to several kinds of knowledge limitations, e.g., limited knowledge of the environment, the collective set of tasks and intentions and possibly capabilities of other agents.

1.2.4 Resources: There are often limited shared resources that each agent can apply to tasks.

1.2.5 Allocation: The appropriateness of a particular agent for a task is a function of how well the agent's skills match the expertise required to do the task, the extent to which its limited knowledge is adequate for the task, its current processing resources (problem solving, communication, time, and workload), and the quality of its communication links with other agents. Not only is the appropriateness of an agent to a task important, but the ability for other agents to recognize the appropriateness of an agent for a task is characteristic of the distributed problem solving situation. Also, an agent

may be unaware of the fact that his expertise is the most appropriate for a particular task.

1.2.6 Information: As in many distributed problem solving environments, the group as a whole may not have all of the information necessary to solve the problem. The dynamic environment provides an opportunity to discover new information and means for relaying that information to the appropriate members of the group.

1.2.7 Structure: The structure of a group often affects its performance. Structure refers to the way in which agents are physically and functionally interconnected. These include communication links, dependence links, role links, power links, etc.

Perhaps the most comprehensive way to describe the situation of distributed problem solving is in terms of what it is *not*. The related field of distributed processing provides a vehicle for doing just that. In distributed processing, multiple computers interact in a pre-determined fashion through the sharing of data. The primary focus lies in an optimal static distribution of subtasks, optimal methods for interconnecting processor nodes, resource allocation, and prevention of deadlock. Complete knowledge of the problem has also been generally assumed, while a major reason for distribution is assumed to be load-balancing. In distributed problem solving, agents must not only share data, but they must also *share the problem solving*. Agents must maintain their focus on solving the problem and possibly developing problem solving methods during the solution process.

Distributed problem solving situations occur in any environment in which the agents that possess the necessary expertise to solve the problem are distributed geographically. In addition, the agents do not have direct access to each other, but must pool their resources and expertise in order to find a solution to the problem. In addition,

the agents are limited in their ability to acquire all of the necessary information about the problem.

1.3 CHALLENGES IN DISTRIBUTED PROBLEM SOLVING

The principal goal of research in the field of distributed problem solving is to gain a fundamental understanding of the structure and dynamics of the problem solving process when it includes outputs from multiple interacting agents (i.e., humans or intelligent machines). Many difficulties and opportunities face distributed problem solvers in several domains. Many of these difficulties are particular to the multiple-agent situation (however spatially distributed) as opposed to the single-agent situation. Following is a discussion of the global problems that generally characterize the distributed problem solving situation.

A dominant challenge in distributed problem solving is the tradeoff between communication and computation. In most distributed problem solving situations, the agents are capable of working alone on subtasks that have been partitioned from the main task (computation) and occasionally interacting with other agents (communication). These interactions generally involve requests for assistance on subtasks or the exchange of results.

A process called "task-sharing" has been proposed [Smith and Davis, 81] for use wherever a task can be decomposed into subtasks. This process allows the agents to divide the workload among themselves. Several difficulties surface when agents must decide which agent(s) will perform which subtask(s). Some of these difficulties include:

- 1). assigning subtasks to the agents as part of the group problem solving activity.

In single-agent problem solving, the agent is typically given its task as part of the problem definition.

- 2). partitioning subtasks to the appropriate agents. Because of their incomplete knowledge of the global situation, and the tasks or intentions of other agents, the agents within the group have difficulty mapping subtasks to agents with appropriate, available expertise.
- 3). ensuring coverage among all of the subtasks. All tasks must be assigned to some agent and extra or redundant agents should not be assigned subtasks. For example, in air-traffic control, if the task is to solve a possible spatial conflict, it may be critical to ensure that only one aircraft detours; if two or more adopt that role, they may possibly create a new collision situation.

When subtasks cannot be performed by independent agents working alone or if it is impossible to decompose the task into subtasks, another form of cooperation may be appropriate, "result-sharing" [Smith and Davis, 81]. In this form, the agents periodically report to each other the partial results they have obtained during execution of individual attempts at the solution. It is assumed that individual agents work on portions of the task that have some degree of commonality, such as interpreting data from overlapping portions of an image. In this case, an *a priori* partitioning of the problem exists, agents simply attempt to solve that portion of the problem which is most closely related to their individual area of expertise, or about which they have the greatest amount of incoming data. Therefore, the distributed problem solvers may face severe difficulties in *coordinating task execution*. Like single-agent tasks or subtasks, group tasks may not be independent. Temporal or logical dependencies may exist. In addition, tasks that are not logically connected may interact through shared resources. Single-agent problem solvers may have difficulties in handling nonindependent tasks or subgoals [Sussman, 75], but these difficulties multiply for distributed problem solvers, because of their limited knowledge in many aspects of the situation. If two agents have only local knowledge (e.g., if they know only the local environment and only their own tasks and intentions)

they will not be able to prevent negative interactions between their actions. Similarly, without some knowledge of others' tasks and intentions, positive interactions, the essence of effective distributed problem solving, will be discouraged. In summary, a main challenge in distributed problem solving is that agents must produce results, which are not only appropriate to the relevant subtask, but combine effectively with the solutions of other agents solving dependent tasks.

In the foregoing paragraphs, the motivation for research, the key characteristics and the major challenges of distributed problem solving have been presented. Section 2.0 presents a summary of the major research issues in this field and describes the emphasis of future research relevant to the distributed problem solving situation. Section 3.0 introduces the concept of intelligent communication within a distributed problem solving environment and discusses the associated functional requirements. Specifically, several different communication protocols are defined and tactical command situations that impose different communication requirements are presented. In addition, factors that affect the problem-solving ability of the group are identified along with the knowledge requirements for intelligent communication via an "Intelligent Executive" or mediator are also presented in Section 3.0. The conceptual framework for analyzing the distributed problem solving problem along with the functional architecture and processing requirements of the Intelligent Executive are presented in Section 4.0. Specifically, the concept of operation of the Intelligent Executive is specified in this section. Section 5.0 offers concluding remarks and suggests areas for future research.

2.0 LITERATURE REVIEW

2.1 MAJOR RESEARCH ISSUES

Many work products are identified as dealing with the subject of distributed problem solving; producing a literature that is diverse, but serviceable, as a picture of what can be classified as distributed problem solving research. A minimum definition of a distributed problem solving group is that it must include at least two agents, that these agents must have some degree of information and/or control autonomy, and that some nonempty subset of the agents display some form of intelligence (capability of reasoning, planning, problem solving, etc.). This definition, clearly, gives much leeway as to what may be considered distributed problem solving; research has explored many of the various dimensions of distribution. Some of these will be discussed in the following sections.

Distributed problem solving has been used in two very different senses in the literature, both as "using artificial intelligence (AI) techniques to find solutions for communities of agents" (called "planning for multiple agents") and as "using communities of agents to find solutions" (called "distributed problem solving") [Rosenschein, 82]. Each of these partial definitions has a place in the developing sense of what distributed problem solving is; however, the latter problem is considered the more difficult of the two, and occupies a central place in distributed problem solving research.

Many attempts have been made to understand how groups cooperate effectively by describing in formal, computational terms the actions of each agent as the group achieves a collective goal. Previous work in cognitive science helps little in achieving this understanding. Over the past twenty years, *beginning with the pioneering work of*

Newell and Simon [Newell and Simon, 72], cognitive scientists have learned much about the information processing that underlies the problem solving of individuals. Cognitive psychologists have, for example, carefully studied the way in which people play games, solve mathematical problems, and program computers. In a similar way, researchers in artificial intelligence have developed computational models of how single agents might construct blocks-world artifacts, do medical diagnosis, and plan genetics experiments, to mention only a few. These efforts have resulted in the development of a variety of techniques for modeling the environment, planning under uncertainty, and executing complex sequences of actions. Unfortunately, recent work suggests that the representations of knowledge [Konolige, 81; Appelt, 82] and planning expertise [McArthur and Klahr, 85] required of agents in distributed or group problem solving situations are quite different than those required for single-agent problem solvers. Organizational psychologists have explicitly studied group performance [Dalkey, 77], but because of the difficulties of representing multiple disparate world views and of specifying sequences of activities within and between agents, their theories are usually expressed informally.

Although some of the work in human group problem solving gives some insight into behaviors of groups, very little may be applied to the distributed problem solving situation. The primary difference between the two is that in distributed problem solving, unlike group problem solving, each agent has not only a different view of the problem but has only a *partial* view of the problem. Agents are limited by their "sensing" capabilities and therefore do not behave in the same manner as someone who has a *global* view of the situation.

Investigators have been drawn to distributed problem solving by the immense challenges of the field and the belief in the ultimate benefits to be gained, both intellectually and operationally, from cooperating, distributed agents. There are many unsolved problems in this field, and current work has only begun to address them; this

section outlines some of the key issues that are being faced in distributed problem solving research. The list is not composed of completely distinct items, since there is overlap between some of these research topics. Nevertheless, the enumeration that follows outlines and emphasizes the primary issues, and is a useful means of organizing them for better understanding.

2.1.1 Global coherence out of local actions

One important high-level issue in distributed problem solving is how to achieve global coherence out of the many local actions taken by a distributed group. By global coherence, we mean that the end result of the group's activity should have certain desirable properties; these properties may differ depending on the group, but usually include such things as correctness, or agreement on a "final answer" either by all, most or some significant subset of the agents.

The distributed problem solving requirement of autonomy is crucial; clearly, if all local actions are dictated by a central authority (with its own utility function) then those local actions serve the centrally dictated common good. When local agents have control over their own actions, however, and have their own goals and their own utility functions, then they will pursue their own or the group's interests as they perceive them. The question then arises of how these locally motivated agents can be guided so as to fulfill some global goals. Dictating certain principles to agents is an option, but the issue of which principles of cooperation should be utilized is still unsettled.

2.1.2 Organization paradigms

Closely related to the issue of globally coherent actions is that of organization structures. Deciding how to structure a group of agents (that is, deciding how they will

interact in terms of control and communication) is a crucial first step in achieving coherent interaction. Several organization paradigms have been explored by distributed problem solving researchers in their search for effective control schemes: unstructured communication (i.e., total connectivity among agents) [Hewitt, 80], hierarchical organizations [Corkill and Lesser, 83; Wesson, et.al., 80], selectively connected nets of agents [Lesser and Erman, 80; Christie, 54], etc. The underlying paradigms of interaction within these organizations have also been varied; some researchers have used an economic model of interaction [Doyle, 85; Malone, 82], some have used "master/slave" analogies [Rosenschein, 82], some biological analogies [Kennedy, 62], while yet others have theorized totally cooperative agents [Davis and Smith, 81].

The work on organizational paradigms, like much else in distributed problem solving, remains in its infancy. Examination of this issue, in conjunction with other, related issues such as control, leaves much work for the future.

2.1.3 Distribution of control

Inextricably connected with the question of global coherence is that of control schemes: how should the individual agents in a distributed group be controlled (so as to further some global utility function)? The essence of this question lies in how group control should be distributed. At one end of the spectrum there are groups where all control is exercised by a single authority (though the agent that is "controller" at any given time or for any given task may change) [Rosenschein, 82]. At the opposite end are groups of totally autonomous agents whose only control over one another is through requests that may be refused. In the middle lie systems [Lesser and Corkill, 83] whose agents interact through subtle processes of hypothesis or desire transferral; control is exercised through a weighting of importance between an agent's own desires and/or hypotheses and the desires and/or hypotheses of other agents.

The advantage of centralized control is clear: global considerations, at least to the extent they are seen by the central agent, can be easily and directly translated into pertinent actions, and there is virtually no trouble in achieving global coherence. With distributed control, on the other hand, decisions can be made without the bottleneck of a central decision maker by reducing processing pressure at the center. Local reaction time can be improved and vulnerability of the whole can be reduced. Distributed control may also be more suitable to the specifications of a particular problem or the requirements for the form its solution should take (political or operational). In any case, group control needs to be understood more fully for distributed problem solving groups to operate efficiently; this becomes especially critical in large groups, consisting of more than just a few agents (where simpler control schemes fail to produce satisfactorily efficient behavior). While some preliminary research has been pointed in this direction [Corkill, 80]; much remains to be done.

2.1.4 Incomplete, possibly inconsistent, views of the world

When problem solving investigators began exploring distribution, one of the first new problems they encountered was that of differing world views among agents. Before, research had dealt with a single agent whose view of the world was often assumed to be correct. The move to distribution argued against continuing to work with this simplifying assumption, whose plausibility, in any case, had never been very high. In particular, when agents are independently performing actions, it is unreasonable (especially in large groups) to update every agent's view of the situation after each action. Thus, with several agents it seems realistic to expect variations among the world models of the agents; at the very least, we might expect gaps in the knowledge of agents, where agent A knows something that agent B doesn't know. A more extreme assumption is that agents have conflicting beliefs.

Research has focused on how agents can communicate knowledge about the world to each other; a less well-developed area of investigation has been belief revision, how beliefs that are discovered to be incorrect can be corrected by a better informed agent. In cases of both "knowledge gaps" and "incorrect knowledge," coherent action will often require agents to have overlapping models of relevant parts of the world. Lacking harmonious beliefs as to the state of this model, two agents will rarely be able to cooperate in achieving a particular goal state.

2.1.5 Communication vs. Computation

There is a tradeoff to be made between communication and computation in distributed groups. It will often be the case that certain information is needed (about the world or about another agent) and could be gained in one of two ways: through querying of an agent who already has, or can more readily obtain, the information; or by performing the necessary computation to answer the question oneself (when this is possible). Similarly, if an agent has a rich enough model of another agent's beliefs and reasoning processes, he could deduce that agent's actions through computation in place of direct communication and querying. The circumstances under which computation should be used in place of communication is an open research topic. Relevant considerations include the relative costs of the two options, limitations on the communication bandwidth, limitations on effective processing time, concerns about communication secrecy (e.g., wanting to limit communication in certain circumstances), and expectations of either method's successfully providing the information desired.

In distributed problem solving groups, the assumption of limited bandwidth (setting maximum limits on the amount of communication over time) is almost universal. This is because limited bandwidths dominate groups in the real world, and also because it is a powerful motivating assumption (guaranteeing autonomy among agents, a "loosely-

coupled" system). Communication is assumed to be limited and more costly than computation using some metric of cost. Without this supposition, the system becomes strongly linked and more like a distributed processing environment, rather than a collection of independent agents.

When a decision is to be made between communication and computation, there is the meta-level issue of how much computation should be expended on making the decision itself [Barnett, 84; Rosenschein and Singh, 83]. In any case, very little is understood about how this decision should be made and the heuristics that might intelligently guide the process [Lesser and Erman, 80].

2.1.6 Resource tradeoffs

Related to the notion of communication vs. computation is the more general problem of resource tradeoffs. In distributed problem solving groups, it is often possible to devise several plans to accomplish some goal; each of these plans may have things that recommend its implementation. For example, consider a plan to format a computer file and print it out. If the file is located at site A, and the printing should go on at site B, where should the formatting take place? Considerations include the availability of the necessary formatting software at the two sites, the cost of doing the formatting at each site, the load on each computer, the size of the formatted and unformatted files, the cost of file transmission, etc. It is necessary to consider these aspects when devising a plan, and to trade off some resources in favor of others.

2.1.7 Synchronization of activities

When agents are performing actions independently, there is the danger that those actions might interfere with one another. Even in situations where their actions can be

guaranteed to combine coherently into an overall plan (through the use, for example, of a central planner), there is the danger that actions will be taken in the wrong order once they are parceled out to separate agents. Early research on this topic has been characterized by two distinct approaches: 1). centrally formulating plans so that they contain the requisite synchronization primitives (i.e., synchronizing the actions of agents) [Rosenschein, 82], and 2). modifying the separate, already existing plans of agents so that they do not conflict [Georgeff, 83], focusing on limiting the combinatorial explosion implicit in examining plan interleavings). Work on synchronization in relatively well-developed, though other, possibly superior, approaches could be undertaken in the future.

2.1.8 Task decomposition

Even when a plan has been fully constructed, there remains the question of how it should be distributed through the system, that is, what agents should carry out what parts of the plan. This question has been partly dealt with in analysis of algorithms (articles on scheduling) [Dolev, 80; Graham, 69; Helmbold and Mayr, 84; Mayr, 81], in systems [Ackerman, 82] and in artificial intelligence [Rosenschein and Singh, 83; Singh and Genesereth, 84].

Although the scheduling problem (in all but trivial incarnations) is difficult, powerful and simple heuristics have been found to aid in the distribution decision. Nevertheless, there remain many areas for improvement. Not enough consideration has been given to the varied nature of system agents; some agents are capable of performing certain tasks but not others, or can perform certain tasks very quickly. These considerations need to be taken into account in the distribution heuristics. Currently, a simplifying assumption is often made that all agents are identical in their capabilities. Ideally, the group should

also use information about agents and their capabilities in constructing the plan, and in dividing it into pieces for distribution.

2.1.9 Reliability and redundancy considerations

One of the hopes for distributed problem solving is that it will enable distributed groups to operate with increased reliability. To accomplish this goal, however, there needs to be more understanding of the tradeoffs involved in achieving greater reliability. For example, one possible way of having a distributed problem solving group fail less often would be for all actions to be performed by all agents (where such overlap is possible). While this would increase reliability through increased redundancy, the cost would clearly be prohibitive in real systems. A lesser amount of redundancy is possible (with a correspondingly lesser amount of reliability), of course, but exactly how much reliability is called for in a particular situation, and how is this best achieved? While work in this area is sketchy (and has been pursued on a mainly intuitive basis by computer systems researchers), it could be a rewarding area for future distributed problem solving research.

2.2 CURRENT RESEARCH EMPHASIS

Distributed problem solving researchers have pursued their aim of designing cooperating agents in a variety of ways. This section briefly categorizes some of the approaches taken. Researchers have, for the most part, stayed within the boundaries of some particular category, with cross-over being rare. This, however, may change as the field matures, and scientists consider complementary ways of approaching their research problems.

2.2.1 Experimentation

Several researchers have approached distributed problem solving with an experimental bent, building simulations of distributed cooperative groups and then conducting performance tests to judge how they work under a variety of parameter values. Most notable among these researchers are Lesser, Erman and Corkill [Lesser and Corkill, 83; Lesser and Erman, 80], and Smith [SIGART, 82]. This approach has quite a bit to recommend it, given the still rudimentary state of knowledge about distributed problem solving. Hard data is badly needed. However, some theoretical foundation is needed, and though much illuminating research has been done through distributed problem solving experiments, the work suffers from a relative abundance of scattered data and few organizing principles. Fortunately, these principles are beginning to emerge [Corkill and Lesser, 81; Corkill, Lesser and Hudlicka, 82].

2.2.2 Formalized theories of knowledge and belief

Part of the work in distributed problem solving is currently directed towards developing formal theories of knowledge and belief [Appelt, 82; Konolige, 83]. This work, mainly taking place at SRI, grew out of the (single agent) work in knowledge and belief pioneered by McCarthy and by Moore; the SRI researchers are extending the concepts (and developing new ones) for the multiple-agent domain. Recent work out of Stanford University [Moses, 86] specifically addresses the role of knowledge in a distributed environment. A general framework for defining knowledge in a distributed system is given. These research efforts are of great interest and importance within both the distributed problem solving community and the problem solving community as a whole; a major concern about the approaches being taken, however, has to do with whether the efficiency of their implementations will be sufficient for applications. For example, some

of their representations make use of the semantics of possible worlds, and this has led to inefficient computation. As work continues on improving these representation schemes, progress can also be expected in making them computationally more tractable.

2.2.3 Principles of cooperation

Some distributed problem solving researchers have devoted their time to developing "principles of cooperation," which should guide the actions of cooperative agents [Davis, 81; Wesson et.al., 80]. These principles comprise such rules of thumb as having agents act predictably, and committing themselves as late as possible to any particular restrictive course of action. The work on such principles was a worthwhile first step towards cooperative system design; unfortunately, the groundwork has not been followed by concrete implementations of these principles in working groups. It is difficult to determine if this is because an innate "fuzziness" in the approach precluded successful follow-up work, or whether other considerations applied. In any case, some of these principles, having been explicated, may find their way into future distributed problem solving implementations by the same investigators or others.

3.0 INTELLIGENT EXECUTIVE FUNCTIONAL REQUIREMENTS

3.1 OVERVIEW

Distributed problem solving agents typically operate in an uncertain environment in which they need to rely heavily on their own knowledge, expertise and information gathering abilities. But in order to effectively pursue and accomplish a common goal, the agents need to intelligently communicate their local "news" and coordinate their mutual actions. It is well recognized from distributed problem solving literature that communication and cooperation are the underpinnings that promote group synergy. Cooperation is the means by which a group optimally utilizes and exploits its mutual strengths and circumvents its weaknesses. Communication among distributed agents is a function which allows those agents to cooperate. Although communication between agents is necessary for effective, cooperative, problem solving, it is just another contributing factor to problem solving that may or may not be used effectively. Considerable expertise is required for effective communication. Such expertise is required to explicitly assess the specific benefits, costs and risks of different communication strategies over different problem solving situations.

This report presents a system concept for intelligent communication and cooperation in an Intelligent Executive. The Intelligent Executive is defined as an agent that possesses meta knowledge, i.e., knowledge about the knowledge that each of the agents possesses. Given this knowledge, the Intelligent Executive is able to assess the benefits, costs and risks associated with the different communication strategies within a particular problem solving situation. On this basis, the Intelligent Executive is able to select the appropriate communication strategy for the group. This section presents the relevant communication strategies and the different situations within which the functional requirements are specified, along with the cost and risk factors that may be present in the

environment. The section ends with the specification of the knowledge requirements of an Intelligent Executive.

3.2 COMMUNICATION PROTOCOLS

Decisions about when and how one agent should communicate with another are continuously required throughout the distributed problem solving process. The different forms of communication or communication strategies are technically known as "communication protocols." Communication protocols provide the types and degree of restrictions (if any) on message-passing among intelligent agents. Three communication protocols that may exist within the distributed problem solving situation are discussed below. Table1 summarizes the advantages, disadvantages and underlying assumptions of each.

3.2.1 Broadcast: Broadcast communication is an indiscriminate form of communication. Any agent can transmit information simultaneously to all other agents within the group. Because of the ability to communicate to all agents, this protocol allows the agents to quickly acquire a complete view of the problem and the surrounding environment. In addition, by sharing their knowledge of the situation with others, the agents can rapidly increase the amount of common information within the group. Whenever a group of agents communicate among themselves, the issue of relevancy becomes important. Some of the transmissions, or messages, may only be relevant to a few agents within the group. Therefore, if an agent receives too much irrelevant information, there is a greater chance of confusion or misinterpretation on the part of the agent. Such interference could ultimately inhibit the progress of the group.

Table 1.
Communication Protocols

<div>COMMUNICATION PROTOCOLS</div> <div>CHARACTERISTICS</div>	BROADCAST	SELECTIVE	AUTONOMOUS
	ADVANTAGES <ul style="list-style-type: none"> - quickly acquire complete view of problem and environment - rapidly increase amount of common information among agents - agents receive all relevant information 	<ul style="list-style-type: none"> - slowly acquire more complete view of problem and environment - slowly increase amount of common information among agents - less chance of allowing interception or detection - agents receive less irrelevant information 	<ul style="list-style-type: none"> - no chance of allowing interception or detection - no cost in terms of time and effort - agents do not receive irrelevant information
	DISADVANTAGES <ul style="list-style-type: none"> - agents may receive large amounts of irrelevant information - highest chance of allowing interception or detection 	<ul style="list-style-type: none"> - acquisition of more complete view of problem and environment is slow - increase in amount of common information among agents is slow - some chance of allowing interception or detection - some cost in terms of time and effort - agents do not receive all relevant information - agents receive some irrelevant information 	<ul style="list-style-type: none"> - agents may never acquire complete view of problem or environment - agents may never increase amount of common information among themselves - agents do not receive relevant information
	ASSUMPTIONS <ul style="list-style-type: none"> - agents do not have to be aware of characteristics or capabilities of other agents 	<ul style="list-style-type: none"> - agents must decide to whom to send messages - agents must be aware of who needs what information - agents must be aware of who can supply what information 	<ul style="list-style-type: none"> - agents must rely on their own knowledge, data and information gathering abilities

An advantage, however, is that whenever a message is relevant to an agent, that agent is sure to receive it under the broadcast protocol.

The greatest disadvantage of the broadcast protocol stems from the fact that a copy of the message being transmitted must be sent to each of the other agents within the group, or it is sent in a form which allows anyone with receiving and deciphering capabilities to intercept it. A situation may exist where there is a risk either of the message being intercepted by an "opposing" agent (e.g., a confidential memorandum in a corporate environment) or of giving away the physical location of the transmitting agent (i.e., detection during military communication in enemy territory). In these situations it would be to the group's advantage to limit the number of messages sent within this protocol.

Finally, an underlying assumption within this protocol is that agents do not have to be aware of the characteristics or capabilities (e.g., knowledge, skills, expertise, etc.) of other agents in the group. An agent can be assured that a message which is relevant to a particular agent will be sent to that agent. It is not incumbent upon the sender of the message to determine the most appropriate receptor of that message.

3.2.2 Selective: The selective communication protocol is more conservative in nature. Under this protocol, agents must select one agent to which to transmit a message. Subsequently, one transmission of that message is sent only to the selected agent. Therefore, it is assumed that agents must be aware of which agents need what information and which agents can supply what information. This may introduce a large decision making component into the communication environment. Although the agents may sometimes select an inappropriate agent, they implicitly reduce the amount of irrelevant information that is distributed to other agents. While this may reduce the chance of confusion or misinterpretation, it also reduces the rate at which agents acquire

a more complete view of the problem and the environment. In addition, it reduces the rate at which the amount of common information among agents is increased.

A desirable feature of the selected protocol is that the chance of allowing message interceptions or agent detections may be lower than under the broadcast protocol (all other factors being equal). Therefore, in situations where the threat of interceptions or detection is high, the selective protocol may be preferred over the broadcast.

3.2.3 Autonomous: The most restrictive form of communication to be discussed is the autonomous communication protocol. No messages are transmitted under this protocol. The agents must rely on their own knowledge, skills, data and information gathering abilities. Several of the benefits of this protocol include:

- no chance of allowing interceptions of messages or detections of agents.
- no cost in terms of time in transmission.
- agents do not receive irrelevant information that may cause confusion or misinterpretations.

In situations where the risks of uncontrolled communication are too high to justify other protocols, the autonomous protocol may be selected. Obviously, there are many disadvantages associated with selecting this protocol. Most importantly, the agents probably will not acquire a complete view of the problem or the environment. They probably will not efficiently increase the amount of common information among them. The agents become completely isolated from each other. Their only knowledge of the other agents or of the actions of other agents must be inferred from the changes which take place in the scenario or environment due to the existence of other agents and the results of their actions.

3.3 COMMAND SITUATIONS

Ideally, appropriate communication protocols are selected by weighing the advantages and disadvantages of each as they pertain to a particular problem situation. The problem "situation" is usually a function of the objective or goal of the problem solving group and the cost and risk factors associated with the environment or scenario. Usually the problem solvers must trade off the benefits of one form of communication against the costs and/or risks within the environment, especially near "acceptable" limits. In order to describe the nature of the relationship between communication protocols and the associated costs or risks, this report will present four examples of military command situations. The command situations will provide a contextual backdrop on which to describe the requirements of the Intelligent Executive. The four command situations below were selected for several reasons:

- they correspond to different phases within a military mission.
- each has a different objective within the overall mission.
- the importance of communication differs in each.
- the emphasis of costs and risks associated with communication differ in each.

Each situation is described below by identifying the objective and the importance of communication within the situation. Table 2 summarizes the following discussion.

3.3.1 Mobilization: The objective of the mobilization phase is to organize and deploy the friendly forces. In addition, the commanders must prepare and maintain an overall strategic posture which is based on information concerning the enemy's

Table 2
Command Situations

CHARACTERISTICS COMMAND SIT.	OBJECTIVE	IMPORTANCE OF COMMUNICATION
MOBILIZATION	<ul style="list-style-type: none"> - to organize and deploy friendly forces and maintain strategic posture 	<ul style="list-style-type: none"> - provides commanders with ability to stay aware of other friendly force activities in order to make appropriate deployment decisions
RECONNAISSANCE	<ul style="list-style-type: none"> - to observe situation in enemy territory without being detected and communicating assessment to commanders 	<ul style="list-style-type: none"> - allows units to relay critical information regarding the enemy situation so that commanders can acquire a "complete picture"
COMBAT	<ul style="list-style-type: none"> - to coordinate forces in an effort to defeat the enemy by using tactics and resources effectively 	<ul style="list-style-type: none"> - allows commanders to be aware of friendly force status in order to evaluate and/or modify tactics and resource allocations
WITHDRAWAL	<ul style="list-style-type: none"> - to organize and execute a retreat from battlefield situation while maintaining strategic posture 	<ul style="list-style-type: none"> - may insure coordination of units so as to prevent collapse of defensive/offensive posture

capabilities. Communication plays an important role when several commanders must coordinate their plans in an effort to make appropriate deployment decisions.

3.3.2 Reconnaissance: The primary focus of a reconnaissance mission is to collect information on the enemy and its capabilities without being detected (assuming retaliation is possible). Communication allows the reconnaissance teams to relay their items of information to the command post so that the decisionmakers may acquire a "complete picture" of the battlefield situation.

3.3.3 Combat: In general, the objective of the combat phase of a mission is to coordinate the friendly forces so that they may defeat the enemy by effectively using tactics and resources. Communication is extremely important in this circumstance only when it is timely. Critical information must be passed to each unit quickly so that the commanders may coordinate their actions to succeed.

3.3.4 Withdrawal: The critical factor in the withdrawal phase is to maintain a strong posture. In other words, the commanders must coordinate their plans so that one unit is not left vulnerable to the enemy. Communication may insure this coordination so as to prevent the collapse of the defensive/offensive posture.

3.4 COMMUNICATION FACTORS

Although the overall goal of the distributed group is to have a successful mission, each phase of the mission, or command situation, has a different objective. In addition, the importance of certain cost and/or risk factors associated with the environment or scenario varies within each command situation. Selection of an appropriate

communication protocol, therefore, is largely dependent upon the impact of the cost and risk factors on the ability to achieve the situation objective. For example, if a successful reconnaissance mission depends heavily on avoiding detection by the enemy, then it is safe to assume that a broadcast communication protocol is not an appropriate selection. However, situations are not always as simple as that. Factors such as timeliness of data, accuracy of transmissions, enemy detection and retaliation capabilities, etc. must be considered before a decision on communication protocol can be reached. The following is a discussion of several cost and risk factors which must be considered as part of the overall protocol selection process.

3.4.1 COST FACTORS: Cost factors can be defined as characteristics of the environment or instances that occur in the scenario that, when present, hinder the solution process in some way. These factors may be present throughout the problem solving process or may occur only at certain times during the process. Either way they have a negative effect on the ability of the agents to achieve the mission objective or simultaneously or subsequently achieve some other mission. Cost factors include:

- **Agent Availability/Unavailability:** Whenever the effectiveness of an agent in the group is "lost", information that agent had or had the ability to acquire are lost or non-information capabilities such as firing a weapon are lost. For example, an agent could experience a mechanical or electrical failure which would prevent it from being able to transmit or receive messages or act positively to achieve the goal. In the military context, it is possible that an agent could be destroyed or killed. The degree to which this loss affects the group is related to several factors:

- the relevancy of the lost information or ability with respect to the command situation or scenario.
 - the amount of information or capability lost that is not shared by any of the other agents.
 - the characteristics of the specific communication protocol.
- **Transmission Characteristics:** Always present in any computer-based communication situation are the physical problems associated with transmitting messages. The degree to which these characteristics affect the problem solving process depends upon the severity of each, the characteristics of the situation at hand, and the characteristics of the specific communication protocol. These problems usually include:
- transmission time; amount of time needed to send a message.
 - delay in receiving a message; the capabilities of the receptor facility or channel could delay the reception of a message.
 - accuracy of each transmission; errors or other events may occur in transmission which could possibly change the meaning of the message or simply make it unintelligible.
- **Timeliness of Data:** Unlike the physical aspect of transmission characteristics, timeliness of data refers to the currency of the information either being sent or received. Data can be very accurate and can be transmitted efficiently through communication channels, but if it is not current and relevant to the situation at that moment, it is untimely. The degree to which it affects the problem solving process depends upon the relevancy of the data with respect to the command situation and the scenario and the characteristics of the specific

communication protocol. Figure 2 summarizes the cost factors which must be considered in assessing the environmental conditions of the situation.

3.4.2 RISK FACTORS: Risk factors can be defined as deleterious characteristics of the environment or instances in the scenario that have a certain probability of occurrence. The associated consequences of these factors are considered to be costly but the occurrence of the factor itself is not deterministic. For example, in the military environment, a threat of detection can exist. The risk associated with that threat is a function of three conditions:

- friendly collector, or sensor, characteristics; passive or "invisible" sensors may reduce the risk of detection.
- enemy detection capabilities; the more sophisticated their capabilities, the higher the probability of detection.
- characteristics of the specific communication protocol.

Message interception, another prevalent communication risk within distributed problem solving situations, is also a function of three conditions:

- friendly message protection capabilities (e.g., cryptic codes).
- enemy interception and interpretation capabilities (e.g., intelligence capabilities).
- characteristics of the specific communication protocol.

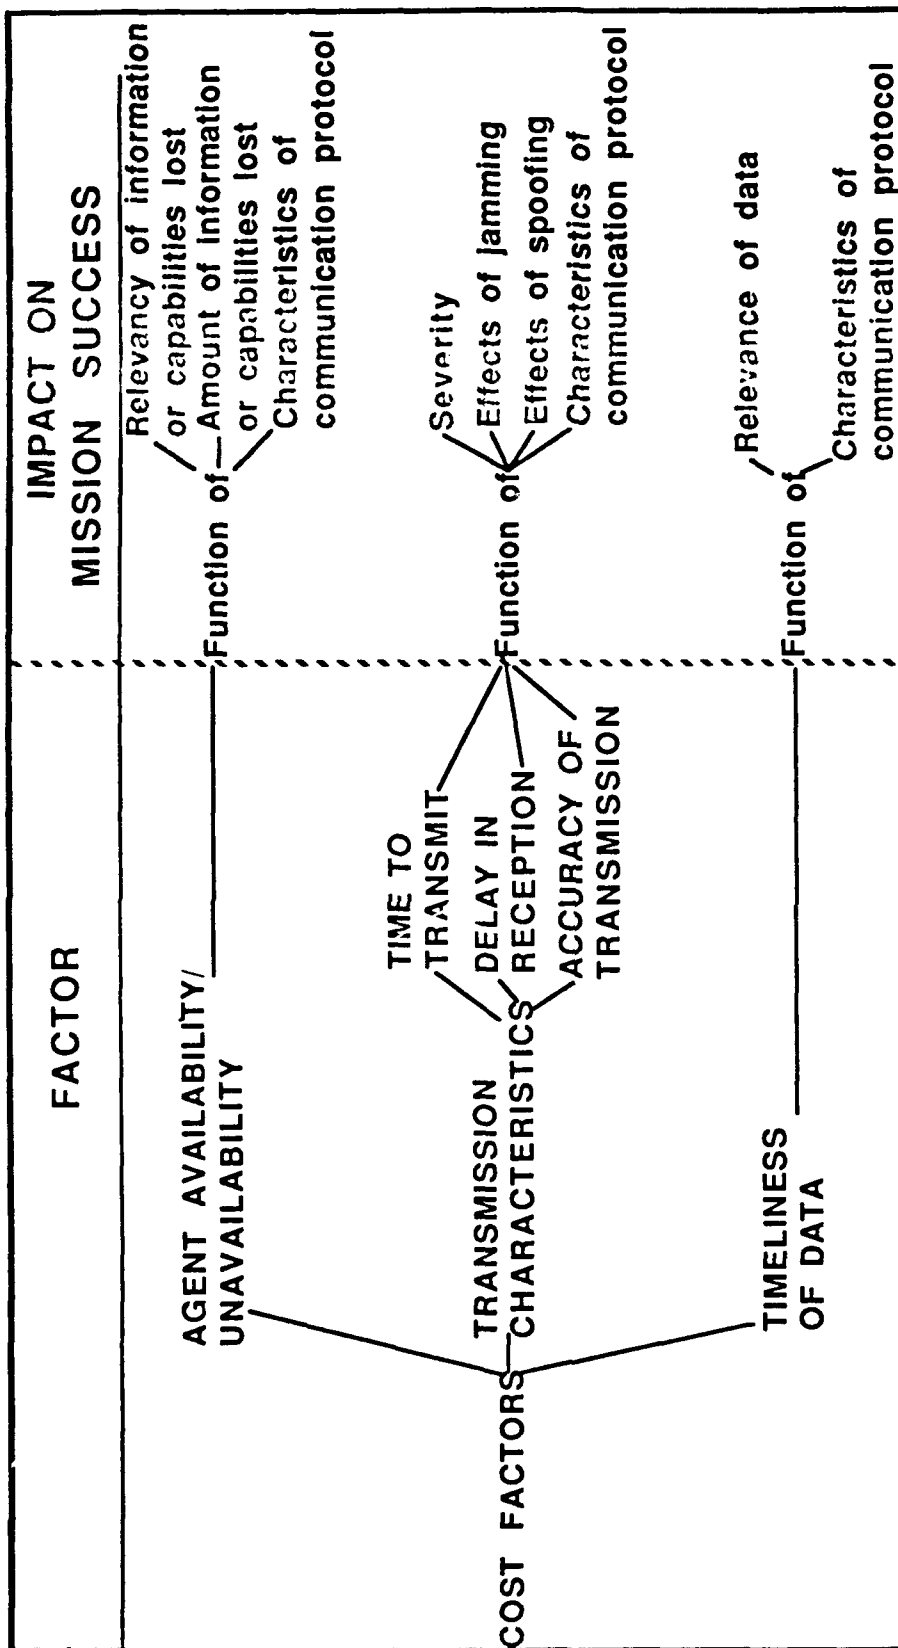


Figure 2.
Cost Factors

The degree to which these risks affect the problem solving process depends upon the consequences associated with the risks. In addition, it depends upon the selection of the communication protocol and how it may increase or decrease the risk within the command situation. Figure 3 shows the risk factors which could be associated with a distributed problem solving situation.

3.5 KNOWLEDGE REQUIREMENTS

During the problem solving process, the Intelligent Executive is responsible for 1) monitoring several aspects of the environment and scenario, 2) maintaining a knowledge base of the problem objective, and 3) maintaining and updating a knowledge base of the agents' informational needs. By doing so, the Intelligent Executive can achieve its goal of assigning the appropriate communication protocol to the group at any point in time. However, there is an assumption that the Intelligent Executive possesses the knowledge and resources required to execute a process for achieving this goal. The specific knowledge required by the Intelligent Executive will be described within the distributed problem solving paradigm presented in Section 1.0

The first of two necessary types of knowledge which must be embodied within the Intelligent Executive is technically known as "domain knowledge." Domain knowledge is the knowledge required of the domain-specific aspects of the problem solving situation. Figure 4 is an amplification of Figure 1 in Section 1.0 which describes the overall distributed problem solving situation. However, Figure 4 contains additional characteristics of the distributed problem solving situation which are relevant to defining the domain knowledge requirements of the Intelligent Executive. Each of the six drawings in Figure 4 highlights a portion of the knowledge bases of one or more of the




Figure 3.
Risk Factors




Figure 4.
Knowledge Base Elements to be Shared Among Agents

agents (knowledge sources). These portions of the knowledge bases must be shared with one or more other agents in order to increase the shared knowledge space (**R**) common to all agents. The Intelligent Executive must somehow be aware of, or estimate the differences between, these knowledge portions.

The ability of the Intelligent Executive to be aware of the different aspects of the domain knowledge is embodied in the *interpretation* the Intelligent Executive gives this knowledge. Figure 5 represents the types of domain knowledge contained within the Intelligent Executive. The difference between Figure 5 and Figure 1 is that the knowledge components of Figure 5 are high-level abstractions of the knowledge components of Figure 1. For example, the knowledge contained within knowledge sources **A**, **B** and **C** of Figure 1 is abstracted in Figure 5 so that the Intelligent Executive knows the type of knowledge embodied in **A**, **B** and **C** but not necessarily the specific knowledge. The body of knowledge, **R'**, that must be shared by the agents to be able to reach a solution is abstracted so that the Intelligent Executive knows what aspects of certain types of knowledge must be common among the agents. The same argument that holds for **R'** is true for the body of knowledge labeled **R**. It represents the aspects of knowledge that the agents currently have in common. Finally, the knowledge specified by **a**, **b**, **c**, **d**, **e**, and **f** in Figure 4 are abstracted in Figure 5 as the aspects of relevant knowledge which must be communicated among the agents in order for the amount of shared knowledge, **R**, to increase.

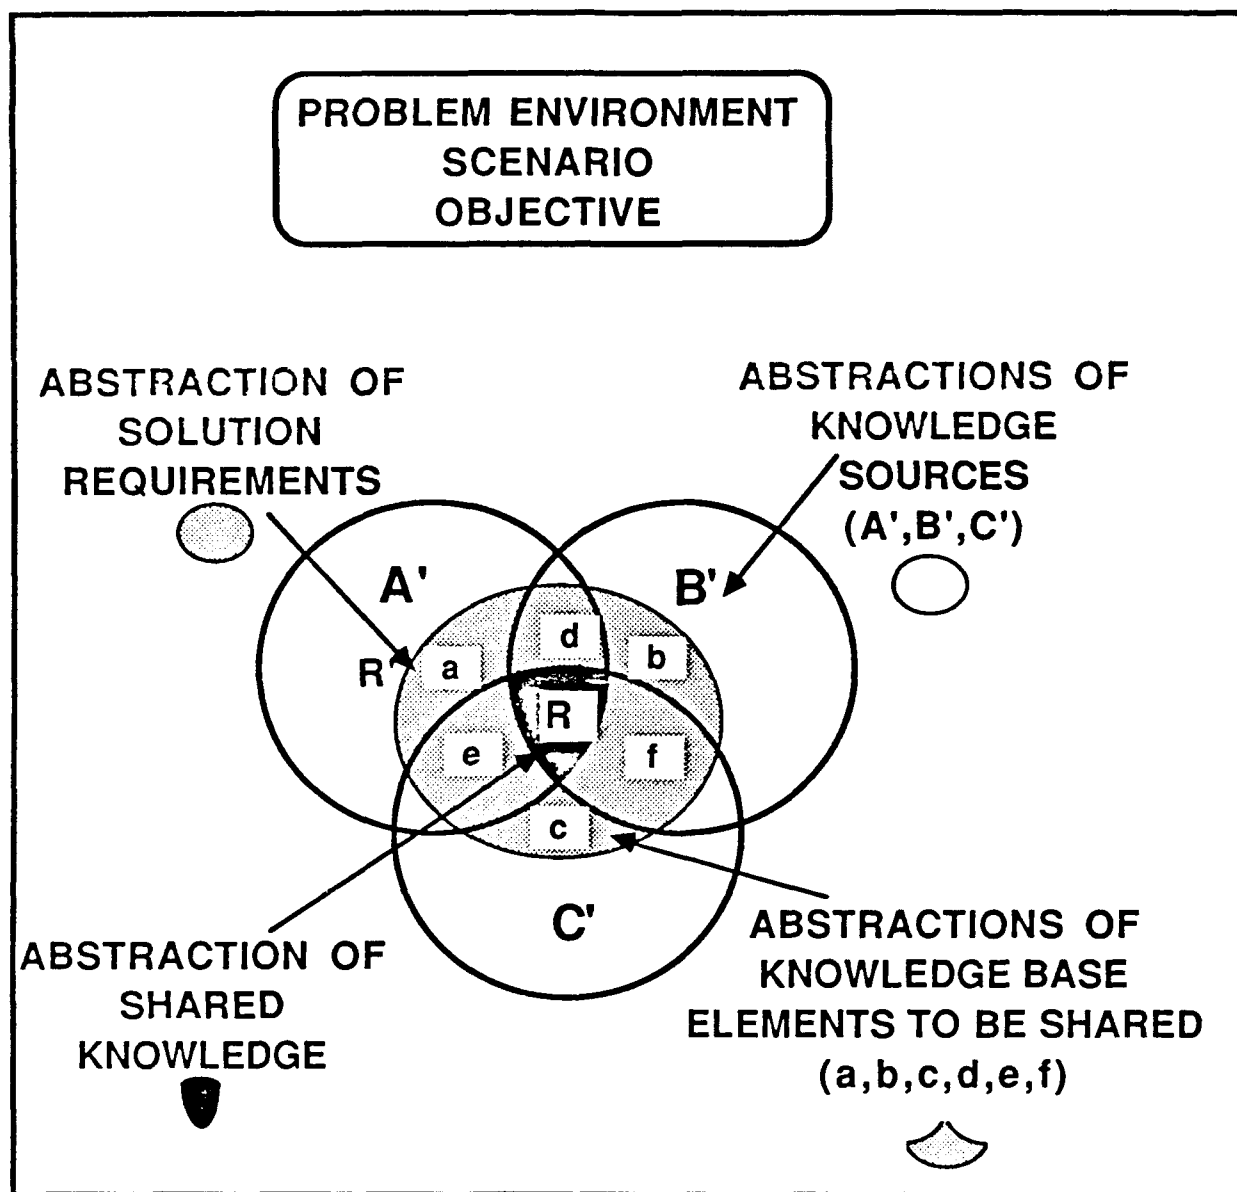


Figure 5.
Intelligent Executive Interpretation of Domain Knowledge

The definitions above represent the intelligent agents and their knowledge relationships with respect to one another. However, these relationships do not exist in vacuum. The contextual backdrop in which they communicate and cooperate is also represented in Figure 5. It contains the bounded problem environment, the scenario and its relevant entities, and the problem solving objective. These are the components of the situation that not only drive the need for the agents to communicate but circumscribe the requirements of the knowledge components defined as **R**, **R'**, **a**, **b**, **c**, **d**, **e**, and **f**. The definition of these components depends on the contextual backdrop in which the problem solving process takes place. Therefore, the Intelligent Executive must keep up with the environmental components to maintain its knowledge base.

The second type of knowledge that must be embodied within the Intelligent Executive is "control knowledge". It is this knowledge that endows the Intelligent Executive with the capability to assess the problem solving situation (domain information) and determine the appropriate *communication strategy or protocol* for that situation. Table 3 summarizes the domain and control knowledge which must exist within the Intelligent Executive.

This section has presented the functional requirements of the Intelligent Executive in terms of the communication protocols it must be able to specify and enforce, the tactical command situation within which it must operate, the communication factors, both risk and cost, that influence the manner in which the Intelligent Executive interprets the environment and the knowledge required for the Intelligent Executive to effectively operate in a distributed problem solving environment.

Table 3.
Intelligent Executive Domain and Control Knowledge

DOMAIN KNOWLEDGE REQUIREMENTS
<ul style="list-style-type: none"> • <u>Types</u> of knowledge embodied in intelligent agents • <u>Aspects</u> of common knowledge among agents required to obtain a solution • <u>Aspects</u> of common knowledge that exists among agents • <u>Aspects</u> of relevant knowledge which must be communicated • Bounded problem environment • Scenario and its relevant entities • Problem solving objective
CONTROL KNOWLEDGE REQUIREMENTS
<ul style="list-style-type: none"> • Criteria for determining the appropriate communication protocol at any point during the solution process

4.0 INTELLIGENT EXECUTIVE CONCEPT OF OPERATION

4.1 OVERVIEW

As a specifier of communication strategies, the Intelligent Executive must be able to assess the cost and risk factors associated with the environment and determine how their severity varies with different communication protocols. Therefore, the Intelligent Executive must have a framework for assessing these factors and their importance within the command situation. The framework must provide criteria for selecting the appropriate communication protocol for the group.

The following is a framework for allowing the Intelligent Executive to assess the appropriate communication protocol for a group of agents during the solution process. Specifically, the framework allows the Intelligent Executive to minimize an objective function which incorporates the effects of the command situation, the scenario, the environment, the cost factors and the risk factors associated with the overall mission.

4.2 CONCEPTUAL FRAMEWORK

It can be assumed, based on the military contextual backdrop, that in order to assess, plan or act effectively the agents must collect information about the environment as accurately as possible. In addition, while pursuing this objective, it is required that the agents minimize the costs and risks associated with their pursuit. For the purpose of defining a conceptual framework in which the Intelligent Executive operates, a military command situation is presented by defining the objective of the mission to be situation

assessment. In addition, the only risk which will be associated with the environment will be the probability of detection by the enemy.

Therefore, the objective function, **F**, can be described as:

$$F = \min f[E(d), E(e)]$$

where

E(d) = expected cost of detection

E(e) = expected cost of error in situation assessment

The elements which describe the objective function (e.g., **E(d)** and **E(e)**) represent the elements of the distributed problem solvers' tradeoff condition. The expected cost of error in situation assessment, **E(e)**, represents the struggle of the problem solver to overcome environmental cost factors which inhibit the flow of information and, therefore, the problem solving process. The expected cost of detection, **E(d)**, represents the environmental risk associated with the attempts to acquire the information necessary to assess the situation. The specific definitions of each of these elements will be stated below.

It is assumed, for the sake of simplicity in specification, that the function describing the relationship between **E(d)** and **E(e)** is additive (although future experimentation may very well prove otherwise). Therefore:

$$F = \min [aE(d) + bE(e)]$$

where

a = f(importance of **E(d)** within the command situation)

b = f(importance of **E(e)** within the command situation)

The representation of the objective function as shown above provides a mechanism for specifying the relative importance of the need for information and the risk associated with trying to obtain it. In some situations, the value of **a** will be larger than the value of **b**. This implies that the importance of the risk associated with communication is greater than the importance of need for accurate information.

The expected costs of both detection and error in situation assessment can be defined in terms of the command situation and the communication factors discussed earlier. Therefore:

$$E(d) = C(d) P(d)$$

where

C(d) = cost of detection

= f(command situation)

P(d) = probability of detection

= f(sensor environment, both friendly and enemy, characteristics of the communication protocol)

and

$$E(e) = C(e) P(e)$$

where

C(e) = cost of error in situation assessment

= f(command situation)

P(e) = probability of error in situation assessment

= f(agent availability/unavailability, transmission, characteristics, timeliness of data, agent stress level, **I_t**)

Table 4 shows the respective descriptions of **C(d)** and **C(e)** for each command situation.

Table 4.
Costs per Command Situation

COSTS COMMAND SIT.	COST OF DETECTION C(d)	COST OF ERROR IN SITUATION ASSESSMENT C(e)
MOBILIZATION	<ul style="list-style-type: none"> - enemy may intercept forces - enemy may interpret strategic plan and ultimately defeat friendly forces 	<ul style="list-style-type: none"> - commanders could generate ineffective plans - commanders may deploy forces in a manner incongruous with strategic plan
RECONNAISSANCE	<ul style="list-style-type: none"> - mission may fail - friendly forces may be captured - counterattack or pre-emptive strike may be triggered 	<ul style="list-style-type: none"> - reports to commanders may be inaccurate - incomplete/erroneous reports could cause strategic plan to be ineffective
COMBAT	<ul style="list-style-type: none"> - if an ambush-type attack, friendly forces may be surprised and defeated - friendly troops could be sized and outnumbered by enemy 	<ul style="list-style-type: none"> - allocation of resources may be suboptimal causing defeat or heavier losses than necessary - commanders may initiate ineffective tactics
WITHDRAWAL	<ul style="list-style-type: none"> - enemy may detect weaknesses in friendly force structure - enemy may intercept forces and prevent withdrawal 	<ul style="list-style-type: none"> - strategic posture may not be maintained - friendly forces may develop weak links

The term I_t in the specification of $P(e)$ refers to the degree of information overlap between agents in the group at any point in time t . During each solution process, there exists a backdrop of shared information between the agents. Ordinarily, the more relevant information the agents have in common, the easier it is for them to communicate with each other. But this common information increases the ability to reach a solution quickly only if its "size" is significant with respect to the amount of common information necessary to reach a solution to the problem. Referring to Figure 6, this shared information can be described as the ratio between the amount of common information between the agents at any point in time (R) and the total amount of common information which the agents must share in order to be able to reach a solution (R').

$$I_t = \frac{R}{R'} = \frac{\text{amount of common information at time } t}{\text{total amount of common information required to be able to reach solution}}$$

As the value of I_t increases, the value of $P(e)$ would tend to decrease (all other factors being equal). One assumption has been made in the design of the framework in which the Intelligent Executive's operates. It is embodied in the fact that no one agent "outranks" any other agent within the group. Agents are allowed to communicate with any agent while keeping with the specified communication protocol.




Figure 6.
Ratio of Shared Information

The framework presented above describes the concept of an Intelligent Executive. By assessing the critical elements of a problem situation which affect the selection of a communication protocol, the Intelligent Executive is able to facilitate the process of distributed problem solving. The following section will present the concept of operation of the Intelligent Executive that would be used to ultimately select the appropriate communication strategy for the group.

4.3 CONCEPT OF OPERATION

The concept of operation consists of seven steps which will be described below. It is by this process that the Intelligent Executive will behave in its role of communication protocol selector. As the distributed group of agents proceed with the problem solving process, the Intelligent Executive must invoke this algorithm to assess the communication requirements of the group. The Intelligent Executive begins this process from the start of the problem scenario, performs it continuously throughout the problem solving process, and concludes it when the intelligent agents have reached a solution to the problem or when the problem ceases to exist.

A concept of operation that is based on the conceptual framework just presented will provide the mechanism for allowing the Intelligent Executive to achieve this objective. This approach is described in Figure 7. The process, which consists of seven separate steps, is outlined below.

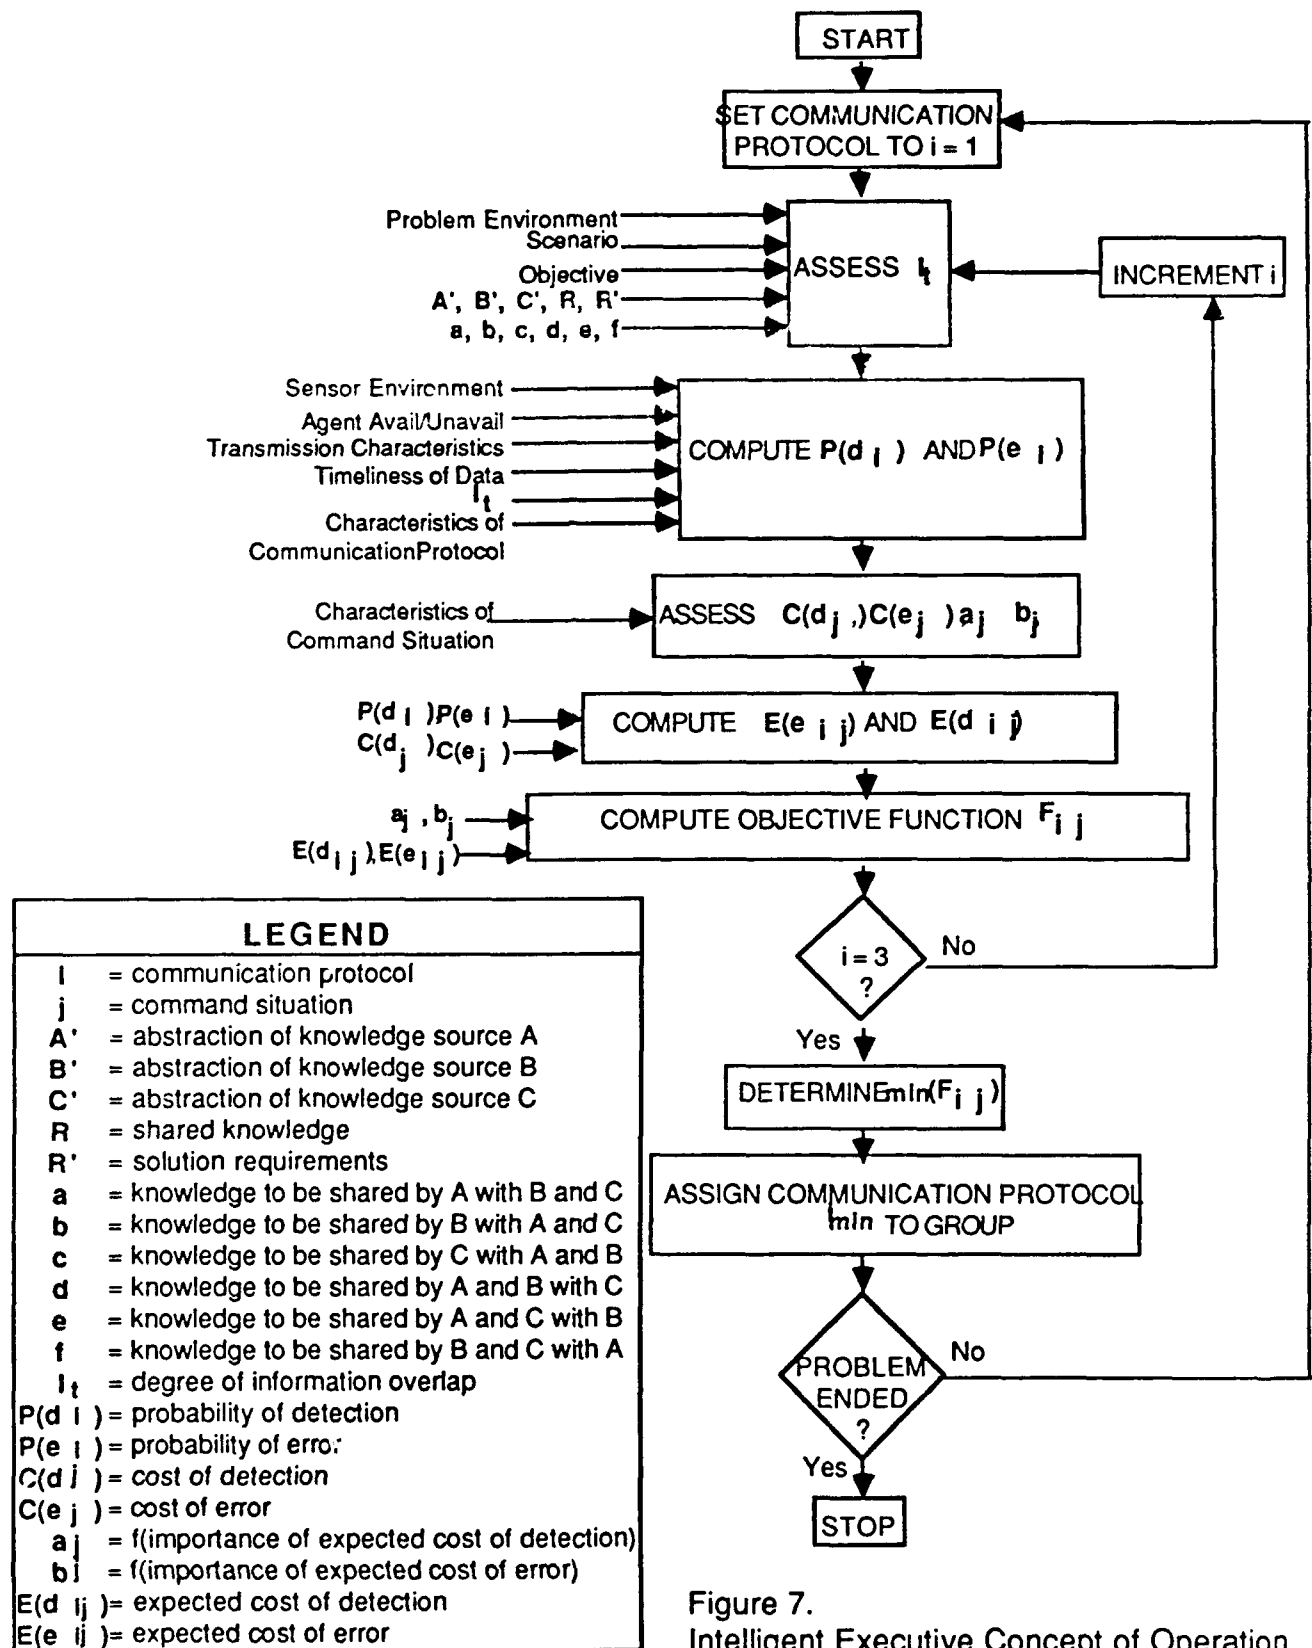


Figure 7.
Intelligent Executive Concept of Operation

Step 1: Assess I_t .

The Intelligent Executive must assess the ratio of R (the amount of common information shared by the agents at time t) to R' (the total amount of common information that must be shared in order for the agents to be able to reach a solution).

Step 2: For each communication protocol i , compute $P(d_i)$ and $P(e_i)$.

The values of the probability of detection, $P(d_i)$, are to be calculated for each communication protocol. Aspects of the environment which contribute to the assessment are the characteristics of the friendly sensors, enemy detection capabilities, and communication protocol characteristics. Calculation of the values of the probability of error, $P(e_i)$, is a function of agent availability/unavailability, transmission characteristics, timeliness of data, the value of I_t , and the communication protocol characteristics.

Step 3: For the current command situation j , assess $C(d_j)$, $C(e_j)$, a_j and b_j .

The costs of being detected, $C(d_j)$, and of making an error in situation assessment, $C(e_j)$, are both functions of the command situation at hand. The relative importance of detection and situation assessment (the values of a_j and b_j) are also functions of the current command situation.

Step 4: Compute $E(e_{ij})$ and $E(d_{ij})$.

The expected costs of error in situation assessment and detection are computed based on the attendant values of probability and costs assessed in the previous steps.

$$E(e_{ij}) = C(e_j) P(e_i)$$

$$E(d_{ij}) = C(d_j) P(d_i)$$

Step 5: Compute the objective function, F_{ij} .

The objective function is computed for each communication protocol i and the current command situation j .

$$F_{ij} = [a_j E(d_{ij}) + b_j E(e_{ij})]$$

Step 6 : Determine $\min (F_{ij}) = F_{i_{\min},j}$.


The three values of F_{ij} are compared and the minimum value is selected. This value ensures that the communication protocol selected (i_{\min}) for the group satisfies the objective that the detection risk and the error costs are kept at a minimum.

Step 7: Assign the communication protocol i_{\min} to the group of agents.

At this point, the Intelligent Executive will dictate the most appropriate communication protocol for the agents to follow which will facilitate their problem solving process in the most efficient manner.

4.4 SELECTION ENVIRONMENT

By following the seven step process described above, the Intelligent Executive would produce results similar to that in Figure 8 for the current command situation (j). The objective function values (F_{ij}) for each communication protocol (i) would be compared and the Intelligent Executive would select the appropriate communication strategy. In the case of Figure 8, the Selective communication protocol would be initiated



Assumptions
for command situation J
$a_j = 1$ and $b_j = 1$ for
$F_{IJ} = [a_j E(d_{IJ}) + b_j E(e_{IJ})]$

Legend
I = communication protocols: Broadcast (B), Selective (S), Autonomous (A)
J = comm. and situation
a_j = f (importance of expected cost of detection)
b_j = f (importance of expected cost of error)
$E(d_{IJ})$ = expected cost of detection
$E(e_{IJ})$ = expected cost of error
F_{IJ} = objective function

Figure 8.
Intelligent Executive Computational Results

with a group of agents. However, circumstances surrounding the problem situation may impose constraints on the selection process such that the minimum value of F_{ij} is not necessarily chosen. For instance, there may exist a maximum acceptable level for the value of $E(d_{ij})$, in which case, the communication protocol selected may not offer the most desirable environment but will satisfy all of the external constraints. Figure 9 shows how this may happen. Although the Selective communication protocol offers the lowest value of the objective function, the value of $E(d_{sj})$ is above the maximum acceptable level. Therefore, the Autonomous protocol is the only solution to the selection process. Figure 10 describes a similar situation in which $E(e_{ij})$ is constrained with a maximum allowable value. In this case, the only viable solution is the selection of the Broadcast protocol which satisfies the constraint.

In essence, there are three solution possibilities which can exist. They are described below.

Case 1: Select $\min(F_{ij})$.

In this case, there are no external constraints placed on any of the values leading to the solution. Therefore, it is simply a case of selecting the minimum value of F_{ij} .

Case 2: Select $\min(F_{ij})$ such that $E(d_{ij}) < E(d_{ij})_{\max}$.

External constraints are placed on the value of $E(d_{ij})$. Therefore, any value of F_{ij} which is comprised of a value $E(d_{ij}) > E(d_{ij})_{\max}$ is ineligible for consideration. The minimum value of F_{ij} of the remaining candidates is then selected and assigned to the group.

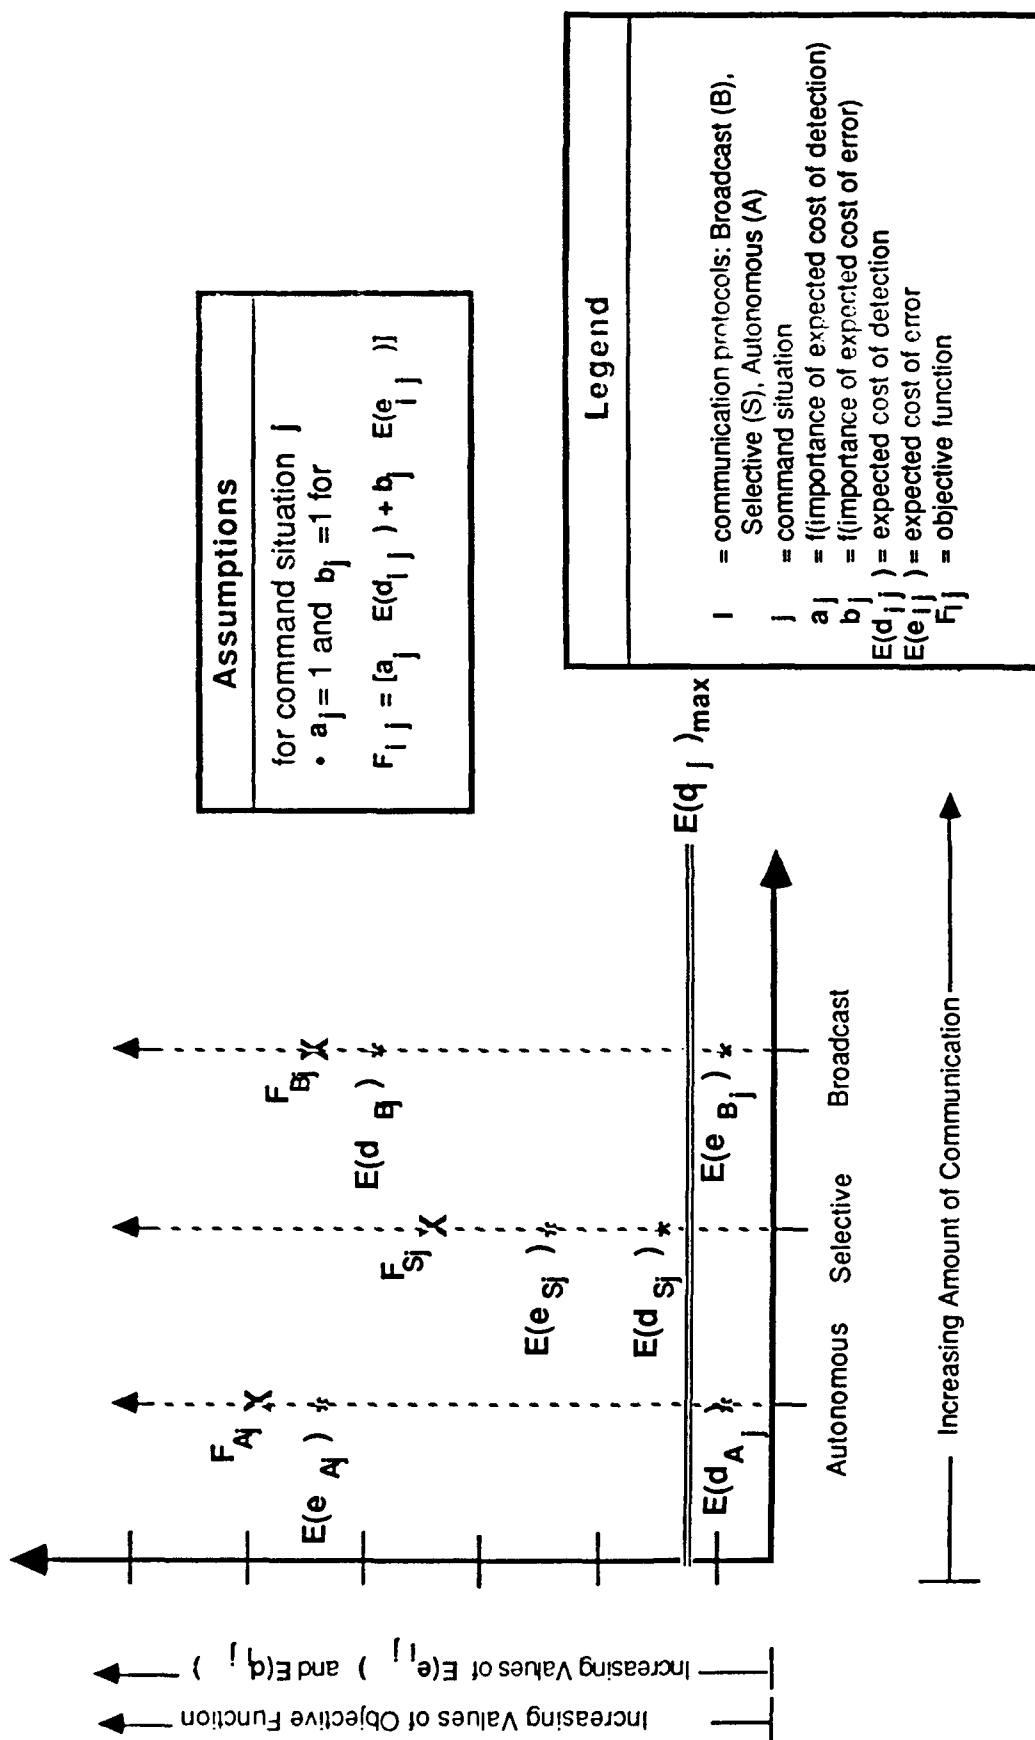



Figure 9.
External Constraints on Expected Cost of Detection



Assumptions	
for command situation	j
$\bullet a_j = 1$ and $b_j = 1$ for	
$F_{ij} = [a_j \ E(d_{ij}) + b_j \ E(e_{ij})]$	

Legend	
i	= communication protocols: Broadcast (B), Selective (S), Autonomous (A)
j	= command situation
a_j	= $f(\text{importance of expected cost of detection})$
b_j	= $f(\text{importance of expected cost of error})$
$E(d_{ij})$	= expected cost of detection
$E(e_{ij})$	= expected cost of error
F	= objective function

Figure 10.
External Constraint on Expected Cost of Error

Case 3: Select $\min (F_{ij})$ such that $E(e_{ij}) < E(e_{ij})_{\max}$.

This situation is similar to Case 2 except that the external constraint is placed on the value of $E(e_{ij})$ instead of $E(d_{ij})$.

Case 4: Select $\min (F_{ij})$ such that $E(d_{ij}) < E(d_{ij})_{\max}$ and $E(e_{ij}) < E(e_{ij})_{\max}$.

In some cases, external constraints prohibit the possibility of a solution at all. For example, if the situations in Figures 9 and 10 existed simultaneously, there would be no communication protocol which would satisfy all of the constraints.

Consequently, the Intelligent Executive would have to rely on a "back-up" strategy.

These may include:

- maintain current communication protocol.
- begin to relax "least important" constraint until a protocol can be selected.
- revert to Autonomous protocol (most conservative strategy).

This section has presented a functional framework and execution process by which the Intelligent Executive will behave in its role as communication protocol selector. This approach is based on an algorithmic method. The inputs through which the algorithm executes include aspects of the problem environment and knowledge characteristics of the intelligent agents within the group. With the specified inputs and the mechanism to execute the process, the Intelligent Executive becomes a fundamental component of the distributed problem solving environment.

5.0 CONCLUSIONS

5.1 OVERVIEW

Communication, both with respect to protocol and content, plays a critical role in distributed problem solving situations. However, a systematic approach for determining the appropriate communication strategy for a particular situation has not been formally put forth by the research community as yet. This report suggests that it is worthwhile to consider a method by which different communication strategies can be practiced or even evaluated within various contexts and under different degrees of risk and cost. Specifically, this report presented the functional requirements of an Intelligent Executive and its concept of operation. The Intelligent Executive presented evaluates a situation and suggests an appropriate communication protocol for the group. The framework presented defines, not only the concept by which the Intelligent Executive performs its functions, but the knowledge that the Intelligent Executive must possess to provide meaningful results.

In an operational context, the Intelligent Executive could serve as a controller of communication. However, the possibilities for use of the Intelligent Executive go well beyond the purely operational environment. A natural setting for this concept is within a testbed environment for the purposes of evaluating distributed problem solving strategies, in any time-varying domain. In each of these settings, the Intelligent Executive could serve as a simulation tool for evaluating agent interactions and communication or novel problem solving strategies. The results of these types of simulations would be extremely valuable to the field of distributed problem solving in that they would guide the development of communication and cooperation strategies in those environments.

Although this report provides a framework for specifying the functionality of an Intelligent Executive, there are many issues which still must be addressed before

implementation of such a concept becomes reality. These issues can help provide guidelines for defining future research. These types of research must be pursued to eventually realize a functioning Intelligent Executive.

In summary, this report presents a framework under which the research issues surrounding the Intelligent Executive concept can be addressed. The following experimental plan describes the objectives, the environment and the variables which could exist in a laboratory setting for the purposes of specifying the functional components of an Intelligent Executive.

5.2 EXPERIMENTAL OBJECTIVE

The experimental framework is designed to facilitate the study of distributed problem solving groups. The behavior of such groups is examined in a simulated, dynamic problem solving environment. The focus of the experimental plan is to provide a controlled environment for addressing issues such as the effects of varying communication protocols on the group's problem solving effectiveness; the effects of varying command situations and their associated cost and risk factors on the probabilities of error and detection/interception; and the evaluation of an objective function for specific problem situations.

5.3 EXPERIMENTAL SET-UP

The proposed experimental set-up is based on a network of personal or micro computers (preferably three agent nodes and one coordinator or experimenter node). The role of the computer in this set-up is primarily to provide a medium for communication among the intelligent agents. However, the computer also allows the experimenter to manipulate the problem environment. In addition, the computer

maintains a record of the agents' actions and provides a facility for viewing the group's solution as it evolves.

The experimental task places a group of three agents in a battlefield situation. The agents are given a "God's-eye view" of a dynamic battlefield environment and are asked to ultimately "assess" the battlefield situation. In order to do so, the agents are given the ability to communicate with each other. They are allowed to transmit messages to each other (either queries or statements) concerning the battlefield developments. However, each agent only sees a "portion" of the overall battlefield. Although their portions may overlap somewhat, no two agents have the same information available to them. In addition, not every portion of the situation is seen by at least one agent. This set-up is characteristic of most distributed problem solving environments. The agents have some common information but as a group they do not always have the "total picture."

The specific task the agents must perform has its roots in situation assessment. The agents are instructed to observe and assess the dynamic battlefield situation. Their task is to arrive, as best as they can, at an interpretation of what is going on in the overall situation. To achieve this goal, the agents must communicate and cooperate with one another in order to better understand the entire environment.

5.4 EXPERIMENTAL VARIABLES

Although the agents are given a facility for communicating, several constraints may be placed on their actions. The variables within this experimental set-up which may be used as constraints on the agents' interactions are described below.

5.4.1 Communication protocols

Communication strategies may be varied in order to observe the effects on end performance or even on the process by which the agents achieve their performance. Several communication protocols such as Broadcast, Selective or Autonomous may be used to constrain/unconstrain the interactions between agents.

5.4.2 Command situations

The command situations or the objective of the situation could be varied in order to study the communication and cooperation strategies employed by the agents in different situations. In addition, with each command situation, the costs of detection/interception and the costs of error in situation assessment may vary. Therefore, the effects of these changes on the agents' actions can be studied.

5.4.3 Probability of detection/interception

The extent to which communication has its consequences is reflected in the probability of detection/interception by the enemy in the scenario. Varying this parameter may change the strategies chosen by the agents and/or their effectiveness in solving the problem.

5.4.4 Time

This variable is usually designed to induce stress within the group performing the task. However, giving the group an indefinite amount of time may also affect the behavior of the agents.

5.4.5 Amount of common information

The "portions" of the battlefield situation which can be viewed by each agent can be varied so that they view more/less of the same information, the group views more/less of the total situation or a combination of the two. The degree to which their portions overlap may affect the agents' problem solving process and/or their end performance.

5.5 SUMMARY

There are many possible combinations of constraints that can be placed on the experimental set-up described above; however, the constraints chosen for any particular experiment are indicative of the results which are desired from it. In other words, the proposed experimental set-up can be viewed as a testbed for defining the functional components of an Intelligent Executive. The testbed could be used to study the sensitivities which exist between the many variables in a distributed problem solving situation.

Finally, it is not suggested that this experimental set-up provides total coverage of all relevant issues to be pursued in the field of distributed problem solving. This set-up is aimed at those issues that are specifically related to the ultimate development of an Intelligent Executive. However, it is suggested that in the process of resolving many of these issues important lessons can be learned to further, not only the insight into the functionality of an Intelligent Executive, but the field of distributed problem solving as a whole.

BIBLIOGRAPHY

- [1] Harold Abelson. Lower bounds on information transfer in distributed computations. *Journal of the Association for Computing Machinery*, 27(2):384-392, April 1980.
- [2] J. M. Abram, C. Y. Chang, V. G. Rutenburg, E. Tse, and R. P. Wishner. *Distributed Decision Making Environment*. Interim Report #RADC-TR-82-310, Advanced Information and Decision Systems, December 1982.
- [3] William B. Ackerman. Data flow languages. *Computer (IEEE Magazine)*, 15(2):15-24, February 1982.
- [4] J. F. Allen. *A Plan-Based Approach to Speech Act Recognition*. PhD thesis, University of Toronto, 1979.
- [5] Douglas E. Appelt. Planning natural-language utterances. In *Proceedings of The National Conference on Artificial Intelligence*, pages 59-62, The American Association for Artificial Intelligence, Pittsburgh, Pennsylvania, August 1982.
- [6] Douglas E. Appelt. A planner for reasoning about knowledge and action. In *Proceedings of the First Annual National Conference on Artificial Intelligence*, pages 131-133, The American Association for Artificial Intelligence, Stanford, California, August 1980.
- [7] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, Inc., New York, 1984.
- [8] Robert Axelrod and William D. Hamilton. The evolution of cooperation. *Science*, 211:1390-1396, March 1981.
- [9] Jeffrey A. Barnett. How much is control knowledge worth?: a primitive example. *Artificial Intelligence*, 22(1):77-89, 1984.
- [10] Avron Barr and Edward A. Feigenbaum (Editors). *The Handbook of Artificial Intelligence*. Volume 2, William Kaufmann, Inc., Los Altos, California, 1982.
- [11] Avron Barr and Edward A. Feigenbaum (Editors). *The Handbook of Artificial Intelligence*. Volume 1, William Kaufmann, Inc., Los Altos, California, 1981.
- [12] M. Benda, V. Jagannathan, and R. Dodhiawala. *On Optimal Cooperation of Knowledge Sources*. Technical Report, Boeing Artificial Intelligence Center, Boeing Computer Services, Bellevue, Washington, August 1985.
- [13] Stephanie Cammarata, David McArthur, and Randall Steeb. Strategies of cooperation in distributed problem solving. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 767-770, The International Joint Conferences on Artificial Intelligence, Karlsruhe, West Germany, August 1983.

- [14] B. Chandrasekaran. Decomposition of domain knowledge into knowledge sources: the MDX approach. In *the Fourth National Conference of Canadian Society for Computational Studies of Intelligence*, Saskatchewan, Canada, May 1982.
- [15] B. Chandrasekaran. Natural and social system metaphors for distributed problem solving: introduction to the issue. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1):1-5, January 1981.
- [16] R. L. Chapman, J. L. Kennedy, A. Newell, and W. C. Beal. The System Research Laboratory's Air Defense experiments. *Management Science*, 5:250-269, 1959.
- [17] Eugene Charniak and Drew McDermott. *Introduction to Artificial Intelligence*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1985.
- [18] Chee Chong and Richard P. Wishner. Distributed hypothesis formation in DSN. In *Proceedings of DARPA Distributed Sensor Network Workshop*, DARPA, 1982.
- [19] Lee S. Christie. Organization and information handling in task groups. *Journal of Operations Research Society of America*, 2:188-196, 1954.
- [20] Paul R. Cohen and Edward A. Feigenbaum (Editors). *The Handbook of Artificial Intelligence*. Volume 3, William Kaufmann, Inc., Los Altos, California, 1982.
- [21] Daniel D. Corkill. *A Framework for Organizational Self-Design in Distributed Problem-Solving Networks*. PhD thesis, University of Massachusetts, Amherst, Massachusetts, 1982. Also published as COINS Technical Report 82-83, Computer and Information Science, University of Massachusetts, Amherst, Massachusetts, December 1982.
- [22] Daniel D. Corkill. *An Organizational Approach to Planning in Distributed Problem Solving Systems*. Technical Report TR 80-13, Department of Computer and Information Science, University of Massachusetts, Amherst, Massachusetts, May 1980.
- [23] Daniel D. Corkill and Victor R. Lesser. The use of meta-level control for coordination in a distributed problem solving network. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 748-756, The International Joint Conferences on Artificial Intelligence, Karlsruhe, West Germany, August 1983.
- [24] Daniel D. Corkill and Victor R. Lesser. *A Goal-Directed Hearsay-II Architecture: Unifying Data-Directed and Goal-Directed Control*. Technical Report TR 81-15, Department of Computer and Information Science, University of Massachusetts, Amherst, Massachusetts, June 1981.
- [25] Daniel D. Corkill, Victor R. Lesser, and Eva Huklicka. Unifying data and goal directed control: an example and experiments. In *Proceedings of The National Conference on Artificial Intelligence*, pages 143-147, The American Association for Artificial Intelligence, Pittsburgh, Pennsylvania, August 1982.

- [26] N. C. Dalkey. *Group Decision Making*. Technical Report UCLA-ENG-7749, School of Applied Science, University of California, Los Angeles, July 1977.
- [27] Randall Davis. *A model for planning in a multi-agent environment: steps toward principles for teamwork*. Working Paper 217, Massachusetts Institute of Technology AI Laboratory, October 1981.
- [28] Randall Davis. Meta-rules: reasoning about rules. *Artificial Intelligence*, 15(3):179-222, December 1980.
- [29] Randall Davis. Content reference: reasoning about rules. *Artificial Intelligence*, 15(3):223-239, December 1980.
- [30] Randall Davis and Reid G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20(1):63-109, 1983.
- [31] Randall Davis and Reid G. Smith. *Negotiation as a Metaphor For Distributed Problem Solving*. MIT AI Memo 624, Massachusetts Institute of Technology, May 1981.
- [32] Gerard Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38:886-893, 1952.
- [33] Danny Dolev. *Scheduling Wide Graphs*. Technical Report No. STAN-CS-80-832, Department of Computer Science, Stanford University, December 1980.
- [34] Jon Doyle. Reasoned assumptions and pareto optimality. In *Proceedings on the Ninth International Joint Conference on Artificial Intelligence*, pages 87-90, The International Joint Conference on Artificial Intelligence, Los Angeles, California, August 1985.
- [35] Edmund H. Durfee, Victor R. Lesser, and Daniel D. Corkill. *Coherent Cooperation Among Communicating Problem Solvers*. Technical Report 85-15, Department of Computer and Information Science, University of Massachusetts, Amherst, Massachusetts, April 1985.
- [36] Edmund H. Durfee, Victor R. Lesser, and Daniel D. Corkill. Increasing coherence in a distributed problem solving network. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 1025-1030, The International Joint Conferences on Artificial Intelligence, Los Angeles, California, August 1985.
- [37] Lee D. Erman and Victor R. Lesser. A multi-level organization for problem solving using many, diverse, cooperating sources of knowledge. In *Proceedings of the Fourth International Joint Conference on Artificial Intelligence*, pages 483-490, USSR, September 1975.
- [38] Richard Fikes. A commitment-based framework for describing informal cooperative work. In *Proceedings of the Cognitive Science Conference*, Berkeley, California, August 1981.

- [39] Richard E. Fikes and Nils J. Nilsson. Strips: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(2):189-208, 1971.
- [40] Mark S. Fox. The intelligent management system: an overview. In H. G. Sol, editor, *Processes and Tools for Decision Support*, North-Holland Publishing Company, 1983. Also published as Technical Report CMU-RI-TR-81-4, Robotics Institute, Carnegie-Mellon University, July 1981.
- [41] Mark S. Fox. Reasoning with incomplete knowledge in a resource-limited environment: integrating reasoning and knowledge acquisition. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 313-318, The International Joint Conferences on Artificial Intelligence, Vancouver, B. C., Canada, August 1981.
- [42] Mark S. Fox. An organizational view of distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1):70-80, January 1981.
- [43] T. Garvey, J. Lowrance, and M. Fischler. An inference technique for integrating knowledge from disparate sources. In *Proceedings of IJCAI*, 7:319-325, 1981.
- [44] Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosenschein. *Solving the Prisoner's Dilemma*. Technical Report HPP-84-41, Heuristic Programming Project, Computer Science Department, Stanford University, November 1984.
- [45] Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosenschein. *Cooperation without Communication*. Technical Report HPP-84-36, Heuristic Programming Project, Computer Science Department, Stanford University, September 1984.
- [46] Michael Georgeff. A theory of action for multi-agent planning. In *Proceedings of the National Conference on Artificial Intelligence*, pages 121-125, The American Association for Artificial Intelligence, Austin, Texas, August 1984.
- [47] Michael Georgeff. Communication and interaction in multi-agent planning. In *Proceedings of the National Conference on Artificial Intelligence*, pages 125-129, The American Association for Artificial Intelligence, Washington, D. C., August 1983.
- [48] J. Goodson, W. Zachary, J. Deimler, J. Stokes, and W. Weiland. Distributed intelligence systems: AI approaches to cooperative man-machine problem solving in C3I. In *Proceedings of the AIAA Computers in Aerospace IV Conference*, pages 1-8, Hartford, Connecticut, October 1983.
- [49] R. L. Graham. Bounds on multi-processing timing anomalies. *SIAM Journal of Applied Mathematics*, 17(2):416-429, March 1969.
- [50] J. S. Greenstein and W. B. Rouse. A model of human decisionmaking in multiple process monitoring situations. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-12(2):182-193, March/April 1982.

- [51] Joseph Y. Halpern and Yoram Moses. *Knowledge and Common knowledge in a Distributed Environment*. Research Report IBM RJ 4421, IBM Research Laboratory, San Jose, California, October 1984. Also published in *Proceedings of the Third Annual ACM Conference on Principles of Distributed Computing*, Vancouver, British Columbia, Canada, 1984.
- [52] P. Haren, B. Neveu, J. P. Giacometti, M. Montalban, and O. Corby. SMECI: cooperating expert systems for civil engineering design. *SIGART Newsletter*, No. 92, April 1985, pp. 67-69.
- [53] B. Hayes-Roth. *A Blackboard Model of Control*. Technical Report HPP-83-38, Heuristic Programming Project, Computer Science Department, Stanford University, August 1984.
- [54] B. Hayes-Roth. *The Blackboard Architecture: A General Framework for Problem Solving?* Technical Report HPP-83-30, Heuristic Programming Project, Computer Science Department, Stanford University, May 1983.
- [55] B. Hayes-Roth and F. Hayes-Roth. A cognitive model of planning. *Cognitive Science*, 3:275-310, 1979.
- [56] B. Hayes-Roth, F. Hayes-Roth, S. Rosenschein, and S. Cammarata. Modeling planning as an incremental, opportunistic process. In *Proceedings of IJCAI-79*, pages 375-383, 1979.
- [57] B. Hayes-Roth and P. W. Thorndyke. *Decisionmaking During the Planning Process*. Technical Report N-1213-ONR, The Rand Corporation, 1980.
- [58] D. O. Hebb. *The Organization of Behavior*. Wiley, New York, 1979.
- [59] J. J. Helly, Jr., W. V. Bates, M. Cutler, and S. Kelem. *A Representational Basis for the Development of a Distributed Expert System for Space Shuttle Flight Control*. NASA Technical Memorandum 58258, May 1984.
- [60] D. Helmbold and E. Mayr. *Fast Scheduling Algorithms on Parallel Computers*. Technical Report No. STAN-CS-84-1025, Department of Computer Science, Stanford University, November 1984.
- [61] Carl Hewitt. The apiary network architecture for knowledge systems. In *Proceedings of the 1980 LISP Conference*, Stanford University, August, 1980.
- [62] Carl Hewitt. Viewing control structures as patterns of passing messages. *Artificial Intelligence*, 8:323-364, August 1977.
- [63] Yu-chi Ho. Team decision theory and information structures. In *Proceedings of the IEEE*, 68(6):pages 644-654, June 1980.
- [64] Douglas R. Hofstadter. Metamagical themas - the calculus of cooperation is tested through a lottery. *Scientific American*, 248(6):14-28, June 1983.

- [65] Douglas R. Hofstadter. Metamagical themas - computer tournaments of the prisoner's dilemma suggest how cooperation evolves. *Scientific American*, 248(5):16-26, May 1983.
- [66] Nigel Howard. The theory of meta-games. *Yearbook of the Society for General Systems Research*, XI:167-186, 1966.
- [67] M. N. Huhns, L. M. Stephens, and R. D. Bonnell. Control and cooperation in distributed expert systems. *IEEE*, pp. 241-245, April 1983.
- [68] S. Kahne, I. Lefkowitz, and C. Rose. Automatic control by distributed intelligence. *Scientific American*, 240(6):78-90, June 1979.
- [69] John L. Kennedy. The system approach: organizational development. *Human Factors*, 4(1):25-52, 1962.
- [70] Kurt Konolige. A deductive model of belief. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 377-381, The International Joint Conferences on Artificial Intelligence, Karlsruhe, West Germany, August 1983.
- [71] Kurt Konolige. A first-order formalization of knowledge and action for a multi-agent planning system. *Machine Intelligence*, 10, 1981.
- [72] Kurt Konolige and Nils J. Nilsson. Multiple-agent planning systems. In *Proceedings of the First Annual National Conference on Artificial Intelligence*, pages 138-142, The American Association for Artificial Intelligence, Stanford, California, August 1980.
- [73] H. J. Leavitt. Some effects of certain communication patterns on group performance. *Journal of Abnormal Social Psychology*, 46:38-50, 1951.
- [74] Douglas B. Lenat. BEINGS: knowlege as interacting experts. In *Advance Papers of the Fourth International Joint Conference on Artificial Intelligence*, pages 126-133, International Joint Conferences on Artificial Intelligence, Tbilisi, Georgia, USSR, September 1975.
- [75] Victor R. Lesser and Daniel D. Corkill. The distributed vehicle monitoring testbed: a tool for investigating distributed problem solving networks. *AI Magazine*, 4(3):15-33, Fall 1983.
- [76] Victor R. Lesser and Daniel D. Corkill. Functionally-accurate, cooperative distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1):81-96, January 1981.
- [77] Victor R. Lesser, Daniel Corkill, J. Pavlin, L. Lefkowitz, Eva Hudlicka, R. Brooks, and S. Reed. A high-level simulation testbed for cooperative distributed problem solving. In *Proceedings of the Third International Conference on Distributed Computer Systems*, October 1982.

- [78] Victor R. Lesser and Lee D. Eрман. Distributed interpretation: a model and experiment. *IEEE Transactions on Computers*, C-29(12):1144-1163, December 1980. Special issue on distributed processing.
- [79] Victor R. Lesser and Lee D. Eрман. A retrospective view of the Hearsay-II architecture. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 790-800, International Joint Conferences on Artificial Intelligence, Massachusetts Institute of Technology, Cambridge, Massachusetts, August 1977.
- [80] Victor R. Lesser, J. Pavlin, and S. Reed. Quantifying and simulating the behavior of knowledge-based interpretation systems. In *Proceedings of the First Annual National Conference on Artificial Intelligence*, pages 111-115, The American Association for Artificial Intelligence, Stanford, California, August 1980.
- [81] Thomas W. Malone. *A Decentralized Method for Assigning Tasks to Processors*. Research Memorandum, Cognitive and Instructional Sciences Group, Xerox Palo Alto Research Center, August 1982.
- [82] Ernst W. Mayr. *Well Structured Parallel Programs Are Not Easier to Schedule*. Technical Report No. STAN-CS-81-880, Department of Computer Science, Stanford University, September 1981.
- [83] D. McArthur and P. Klahr. *The ROSS Language Manual*, Technical Report No. N-1854-1-AF, The Rand Corporation, September 1985.
- [84] D. McArthur, R. Steeb, and S. Cammarata. A framework for distributed problem solving. In *Proceedings of The National Conference on Artificial Intelligence*, pages 181-184, The American Association for Artificial Intelligence, Pittsburgh, Pennsylvania, August 1982.
- [85] B. G. McDaniel. Issues in distributed artificial intelligence. *IEEE First International Conference on Data Engineering*, pages 293-297, 1984.
- [86] Robert C. Moore. *A formal theory of knowledge and action*. Tech Note 320, SRI International, Menlo Park, California, 1984. Also in *Formal Theories of the Commonsense World*, Hobbs, J. R., and Moore, R. C. (Eds.), Ablex Publishing Co. (1985).
- [87] Robert C. Moore. *Reasoning about knowledge and action*. Tech Note 191, SRI International, Menlo Park, California, 1980.
- [88] Yoram Moses. *Knowledge in a Distributed Environment*. PhD dissertation, Stanford University, 1986. Also published as Technical Report No. STAN-CS-86-1120, Department of Computer Science, Stanford University, March 1986.
- [89] A. Newell and H. Simon. *Human Problem Solving*. Prentice-Hall, New York, 1972.

- [90] H. Penny Nii. *An Introduction to Knowledge Engineering, Blackboard Model, and AGE*. Technical Report HPP-80-29, Heuristic Programming Project, Computer Science Department, Stanford University, March 1980.
- [91] Nils J. Nilsson. *Distributed Artificial Intelligence*. Technical Report, SRI International, March 25, 1981.
- [92] Nils J. Nilsson. *Principles of Artificial Intelligence*. Tioga Publishing Company, Palo Alto, California, 1980.
- [93] Nils J. Nilsson. *Research in Distributed Artificial Intelligence*. SRI Research Proposal ECU-79-30, SRI International, Menlo Park, California, April 1979. Part One - Technical Proposal.
- [94] Howard Raiffa. *Decision Analysis, Introductory Lectures on Choices under Uncertainty*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1968.
- [95] S. Randall, S. Cammarata, F. A. Hayes-Roth, P. W. Thorndyke, and R. B. Wesson. *Distributed Intelligence for Air Fleet Control*. Technical Report R-2728-ARPA, The Rand Corporation, October 1981.
- [96] W. R. Reitman. *Cognition and Problem-Solving*. Wiley, New York, 1965.
- [97] Jeffrey S. Rosenschein. *Rational Interaction: Cooperation Among Intelligent Agents*. PhD dissertation, Stanford University, 1985.
- [98] Jeffrey S. Rosenschein. Synchronization of multi-agent plans. In *Proceedings of The National Conference on Artificial Intelligence*, pages 115-119, The American Association for Artificial Intelligence, Pittsburgh, Pennsylvania, August 1982.
- [99] Jeffrey S. Rosenschein and Michael R. Genesereth. *Communication and Cooperation*. Technical Report HPP-84-5, Heuristic Programming Project, Computer Science Department, Stanford University, October 1984.
- [100] Jeffrey S. Rosenschein and Vineet Singh. *The Utility of Meta-level Effort*. Technical Report HPP-83-20, Heuristic Programming Project, Computer Science Department, Stanford University, March 1983.
- [101] S. M. Shatz. Communication mechanisms for programming distributed systems. *IEEE Computer*, pp. 21-28, June 1984.
- [102] SIGART Newsletter, A Quarterly Publication of the Association for Computing Machinery Special Interest Group on Artificial Intelligence. No. 84, April 1983, pp. 3-12.
- [103] SIGART Newsletter, A Quarterly Publication of the Association for Computing Machinery Special Interest Group on Artificial Intelligence. No. 80, April 1982, pp. 13-23.

- [104] SIGART Newsletter, A Quarterly Publication of the Association for Computing Machinery Special Interest Group on Artificial Intelligence. No. 73, October 1980, pp. 42-52.
- [105] H. A. Simon and S. K. Reed. Modeling strategy shifts in a problem-solving task. *Cognitive Psychology*, 8:86-97, 1976.
- [106] Vineet Singh and Michael R. Genesereth. *A Variable Supply Model for Distributing Deductions*. Technical Report HPP-84-14, Heuristic Programming Project, Computer Science Department, Stanford University, May 1984.
- [107] Reid G. Smith. Report on the 1984 Distributed Artificial Intelligence Workshop. *The AI Magazine*, pp. 234-243, Fall 1985.
- [108] Reid G. Smith. *A Framework for Distributed Problem Solving*. UMI Research Press, Ann Arbor, Michigan, 1981.
- [109] Reid G. Smith. The contract net protocol: high-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, C-29(12):1104-1113, December 1980.
- [110] Reid G. Smith. *A Framework for Problem Solving in a Distributed Processing Environment*. PhD thesis, Stanford University, 1978. Also published as STAN-CS-78-700 (HPP-78-28), Department of Computer Science, Stanford University, December 1978.
- [111] Reid G. Smith and Randall Davis. Frameworks for cooperation in distributed problem solving. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1):61-70, January 1981.
- [112] Reid G. Smith and Randall Davis. Distributed problem solving: the contract net approach. In *Proceedings of the 2nd National Conference of the Canadian Society for Computational Studies of Intelligence*, The Canadian Society for Computational Studies of Intelligence, Toronto, Canada, July 1978.
- [113] R. Steeb, S. Cammarata, F Hayes-Roth, and R. Wesson. *Distributed Intelligence for Air Fleet Control*. Technical Report WD-839-ARPA, The Rand Corporation, December 1980.
- [114] R. Steeb, D. McArthur, S. Cammarata, S. Narain, and W. Giarla. *Distributed Problem Solving for Air Fleet Control: Framework and Implementations*. Technical Report N-2139-ARPA, The Rand Corporation, April 1984.
- [115] G. Sussman. *A Computational Model of Skill Acquisition*. American Elsevier, New York, 1975.
- [116] R. R. Tenney and N. R. Sandell. Structures for distributed decisionmaking. *IEEE Transactions in Systems, Man, and Cybernetics*, SMC-11(8):517-527, August 1981.
- [117] R. R. Tenney and N. R. Sandell. Strategies for distributed decisionmaking. *IEEE Transactions in Systems, Man, and Cybernetics*, SMC-11(8):527-538, August 1981.

- [118] P. Thorndyke, D. McArthur, and S. Cammarata. Autopilot: a distributed planner for air fleet control. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 171-177, The International Joint Conferences on Artificial Intelligence, Vancouver, B.C., Canada, August 1981.
- [119] R. Wesson, F. Hayes-Roth, J. Burge, C. Stasz, and C. Sunshine. Network strucutres for distributed situation assessment. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1):5-23, January 1981.
- [120] R. Wesson, F. Hayes-Roth, J. Burge, C. Stasz, and C. Sunshine. *Network Strucutres For Distributed Situation Assessment*. Technical Report R-2560-ARPA, The Rand Corporation, August 1980.
- [121] Patrick H. Winston. *Artificial Intelligence*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1984. Second Edition.
- [122] Ju-Yuan David Yang, Michael N. Huhns, and Larry M. Stephens. An architecture for control and communication in distributed artificial intelligence systems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-15(3):316-326, May/June 1985.
- [123] S. S. Yau and S. M. Shatz. On communication in the design of software components of distributed computer systems. *Distributed Computing Systems Conference*, pages 280-287, August 1982.
- [124] A. Yonezawa and C. Hewitt. Modelling distributed systems. In *Proceedings of the International Joint Conference on Artificial Intelligence*, page 5, 1977.